NASA/CR-97- 205943

NAS2-14283

FINAL
IN-04-CR
OCIT
082349

CSDL-R-2767

FINAL REPORT

EVOLUTION TOWARD A DECENTRALIZED AIR TRAFFIC FLOW MANAGEMENT SYSTEM

November 1996

DRAPER LABORATORY

The Charles Stark Draper Laboratory, Inc.
555 Technology Square, Cambridge, Massachusetts 02139-3563

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

...GENCY USE ONLY (Leave Blank)

2. REPORT DATE
Nov 30, 1996

3. REPORT TYPE AND DATES COVERED
Final (28 Aug 95 - 30 Nov 96)

...ITLE AND SUBTITLE
...volution Toward a Decentralized Air Traffic Flow ...anagement System

5. FUNDING NUMBERS
C NAS2-14283
AFT-1006 01C 806/70110

AUTHOR(S)
Stephan Kolitz
Milton Adams
Mark Abramson
Steven Nelson
William Hall
Ryan Rifkin

PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
C.S. Draper Laboratory
555 Technology Square
Cambridge, MA 02139

8. PERFORMING ORGANIZATION REPORT NUMBERS

SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
NASA Ames Research Center
Moffett Field, CA 94035-1000

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

... SUPPLEMENTARY NOTES

...2a. DISTRIBUTION/AVAILABILITY STATEMENT
A: Approved for public release; distribution unlimited

12b. DISTRIBUTION CODE

3. ABSTRACT *(Maximum 200 words)*

There are three major components of the work reported here. The first is the development of alternative Air Traffic Flow Management (ATFM) concepts ranging from centralized to decentralized. The second is the development of metrics and tools to be employed in the analysis of alternative ATFM concepts that are in this range. The third is the application of these tools and metrics in the analysis of three ATFM operational concepts: passive, current, and collaborative. *Passive* refers to no ATFM and is used as a baseline; *current* is the currently employed approach to ATFM; and *collaborative* refers to a hypothesized ATFM concept that is more decentralized than current ATFM and is a concept in which the FAA allocates a set of arrival slots for each airline and each airline individually determines the assignment of each of its particular aircraft to each slot in its allocated set of slots. Two classes of capacity scenarios are used. The first is a case for which there is VMC throughout the system. The second represents the case of a weather front moving up the east coast, causing IMC for several hours (with uncertain start times) at key airports. Cases with and without en route free flight (User Preferred Routing) and CTAS are modeled. Appropriate metrics (e.g., delay, tardiness) for each of the operational and capacity scenarios have been calculated and are presented.

14. SUBJECT TERMS

15. NUMBER OF PAGES
166

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	UL

Final Report

Evolution Toward a Decentralized Air Traffic Flow Management System

November 30, 1996

Prepared for:
National Aeronautics and Space Administration
Ames Research Center
Flight Management and Human Factors Division, Code AF
Moffett Field, CA 94035-1000

Funding provided by:
NASA Contract NAS2-14283(MJH)

Submitted by:
The Charles Stark Draper Laboratory
555 Technology Square
Cambridge, Massachusetts
02139

DRAPER LABORATORY

SUMMARY

There are three major components of the work reported here. The first is the development of alternative Air Traffic Flow Management (ATFM[1]) concepts ranging from centralized to decentralized. The second is the development of metrics and tools to be employed in the analysis of alternative ATFM concepts that are in this range. The third is the application of these tools and metrics in the analysis of three ATFM operational concepts: passive, current, and collaborative. *Passive* refers to no ATFM and is used as a baseline; *current* is the currently employed approach to ATFM; and *collaborative* refers to a hypothesized ATFM concept that is more decentralized than current ATFM and is a concept in which the FAA allocates a set of arrival slots for each airline and each airline individually determines the assignment of each of its particular aircraft to each slot in its allocated set of slots. The collaborative scenario is one that realistically could be implemented in the medium term; many of its elements are ready to be put in place today. Comparative results for each of these three concepts were generated using January 13, 1993 OAG data as a representative schedule. Two classes of capacity scenarios are used. The first is a case for which there is VMC throughout the system, called the "blue sky" scenario. The second represents the case of a weather front moving up the east coast, causing IMC for several hours (with uncertain start times) at key airports; this scenario is called the "weather front" scenario. Cases with and without en route free flight (User Preferred Routing) and CTAS are modeled. Appropriate metrics (e.g., delay, tardiness) for each of the operational and capacity scenarios have been calculated and are presented.

The collaborative concept investigated here can be viewed as a possible next step in the transition of the ATFM system toward increased decentralized collaborative decision making. The kinds of analyses that the tools we have developed under this effort make possible for this and other anticipated candidate approaches will lead to more systematic and rigorous approaches to the evaluation of the risks and benefits of a spectrum of other proposed concepts. The development of the tools and simulation environment which underlies our ability to perform comparative analyses is considered to be the most significant contribution of our AATT supported research and development effort.[2]

1. INTRODUCTION

ATFM is one of the two major components of Air Traffic Management (ATM), the other being Air Traffic Control (ATC). ATFM includes all activities related to the management of the flow of aircraft and to the management of related system resources from "block to block," including strategic flow management of airport arrival and departure capacities, tactical en route flow management, near terminal area flow management and ground traffic flow management.

ATFM has become increasingly critical to the successful operation of the Air Traffic system both in the United States and in Europe. Continuing growth in system traffic demand is not

1 What we refer to herein as ATFM is explained in Section 2.3.

2 In addition to principal funding provided through NASA Contract NAS2-14283(MJH) under the Advanced Air Transportation Technology Program, this work was partially funded as a Draper Lab Internal Research and Development project. This funding has been used to continue related flow planning research, and to support MIT's role in contributing to the conceptual stages of the project.

being met by corresponding increases in the physical capacity of the system (e.g., new airports) and, therefore must be accommodated by increasing the system's effective[3] capacity through improved management and utilization of the existing system resources. The objective of both strategic and tactical traffic flow management is to match as best[4] as possible the projected demand on the various air transportation system resources (airports, terminal areas, en route sectors) with their available, anticipated capacity. ATFM system functions are most critical to system performance on precisely those days and at precisely those locations where the demand vs. capacity relationship is most unfavorable. The ATFM system relies on a combination of mechanisms, some of which are more global and strategic in nature with longer time horizons (e.g., ground-holding of aircraft prior to departure, ground stop programs and traffic re-routing) and others of which are more local and tactical with shorter time horizons (e.g., miles-in-trail, airborne holding, arrival sequencing and ground traffic management).

The ATFM system is on the verge of a transition that is likely to bring about dramatic changes. This transition is unavoidable, in view of the confluence of several factors including:

(1) The emergence of new technologies that offer the opportunity to correct some of the perceived deficiencies in today's ATFM system.

(2) The expressed preference by the airlines and other aircraft operators for a more decentralized system wherein they participate more broadly in ATFM decision-making.

(3) A general recognition of the need for increased use of decision support tools and automation aids in order to more effectively cope with the highly dynamic environment in which the ATFM system operates, including substantial uncertainty in predictions of demand and available airport capacity when weather conditions deteriorate.

This report explores alternative concepts for modifying the policies and procedures under which the ATFM system operates. These alternative concepts represent stages in the evolution from the current system in which ATFM decision-making is largely centralized within the FAA to a more decentralized approach to decision-making wherein the airlines collaborate in decision-making with the FAA. This evolution is consistent with the decision-making approaches embodied in the "free flight" approach to ATFM. In particular, with more substantial participation from the airlines, air traffic flow will be influenced more directly in ways that accommodate the business objectives of the airlines and the interests of their passengers. At the same time, the FAA will continue to be responsible for the safe operation of the US air transportation system. Thus, a collaboration between the FAA and airlines will be required to insure that system resources will simultaneously address airline business and system safety objectives.

Section 4 elaborates on the themes of centralized and decentralized ATFM and discusses the spectrum of proposed alternative concepts for ATFM ranging from highly centralized (nearly all decision-making is made by the FAA) to highly decentralized (nearly all decision-making is performed by the airlines). Section 5 further explores a viable medium-term partially-decentralized scenario, summarizes metrics that are employed in analyzing and evaluating the

---

3 *Effective* capacity refers to the capacity of a system resource that is realized as a result of the application of a set of policies and procedures for utilizing that system resource.

4 Here, the "best match" is really a multi-objective problem in that "best" is interpreted differently by the various system participants.

various alternatives and describes briefly a simulation testbed that has been developed to generate values for those metrics. In addition to modeling activities of individual aircraft, the simulation testbed must contain behavioral models of the FAA and airlines. Section 5 also raises the fundamental challenge of modeling airline behavior in decentralized ATFM environments and presents an example of such a model, which deals with the preservation of flight bank integrity in hub airports.

2. BACKGROUND

The current Air Traffic Management (ATM) system has served the public well over many years. The traffic has undergone a tremendous increase, while concurrently safety standards have also become considerably more stringent.

The most fundamental shortcoming of the present U.S. airspace system is its limited ability to accommodate uncertainties in capacity (e.g., weather impacts) and in demand (e.g., real-time aircraft requests for more efficient flight path). This system has been in use for approximately 40 years. It was conceived in the infancy of radar and for traffic densities far lower than today's. Over the years, there have been efforts to gain additional capacity to satisfy the rising demand. In the absence of additional automation or new operational concepts, while preserving and improving operational safety, the flexibility to operate efficiently in the national airspace system has been sacrificed. In most instances, the current ATC system dictates the route of flight, altitude and even speed to airline operators. The limitations of this system to accommodate real-time re-planning in response to uncertain capacity and demand result in missed opportunities for economic benefits.

Free flight holds the potential of providing a quantum jump in the air traffic system's operational efficiency, taking advantage of the enabling advances that have been made in communication, navigation and surveillance technologies.

For a new system to be accepted, it has to show significant economic benefits over the existing system, while maintaining a high safety standard. A system-wide analysis is required to capture the intricate relationships among various performance measures.

2.1 CURRENT ATM SYSTEM OPERATIONS

The National Airspace System (NAS) is a collection of interrelated resources supporting aviation in the United States (Figure 2.1-1).

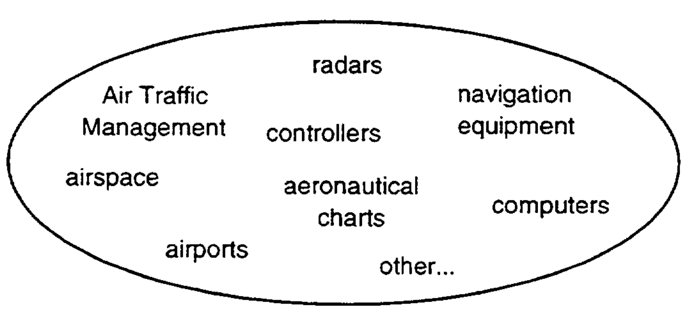

Figure 2.1-1: The National Airspace System

For management purposes, the airspace over the United States is divided into different regions known as sectors. Air traffic moves through sectors along a network of air routes that covers the United States. Air traffic controllers at FAA facilities, monitor collections of sectors on radar display screens and communicate control actions directly to aircraft pilots. This division into sectors is referred to as a *control* division.

At the same time, airspace is classified into regions based on *regulatory* divisions. Airspace regulatory divisions, independent of sector divisions, separate airspace regions into classes which contain specific FAA rules for pilots operating therein, and controllers managing the airspace. Aircraft equipment requirements vary between different classes of airspace. A side view of the airspace is shown in Figure 2.1-2. The entire airspace is divided into sectors, designed to balance controller workload. Control of the various sectors and airspaces are divided among, and sometimes shared by, Airport Traffic Control Towers (ATCTs), Terminal Radar Approach Control Facilities (TRACONs), tower controllers, and Air Route Traffic Control Centers (ARTCC).

Any traffic management scheme must allow for the efficient and safe coexistence of the four types of air traffic: air carrier, air taxi, general aviation, and military aviation. Flights are classified under flight rules (Instrument Flight Rules-IFR or Visual Flight Rules-VFR) according to the scheduled highest altitude and forecast weather conditions. IFR operation requires specific equipment and pilot training. Most air carrier traffic is IFR. IFR traffic receives continuous controller services. On the other hand controllers are not required to provide VFR flights with air traffic and weather advisory services. VFR pilots must provide their own separation distances from all other flights and the terrain. Because of the busy nature of operations in and around airports, all flights (VFR and IFR) are under *positive control* when operating in Class B airspace.

A key aspect of the current ATM is the *flight plan*. The FAA requires that flight plans be filed for all IFR flights. Once a flight plan is submitted to the FAA, the pilot must receive permission to deviate from the originally submitted route. Filing a flight plan is optional for VFR traffic.

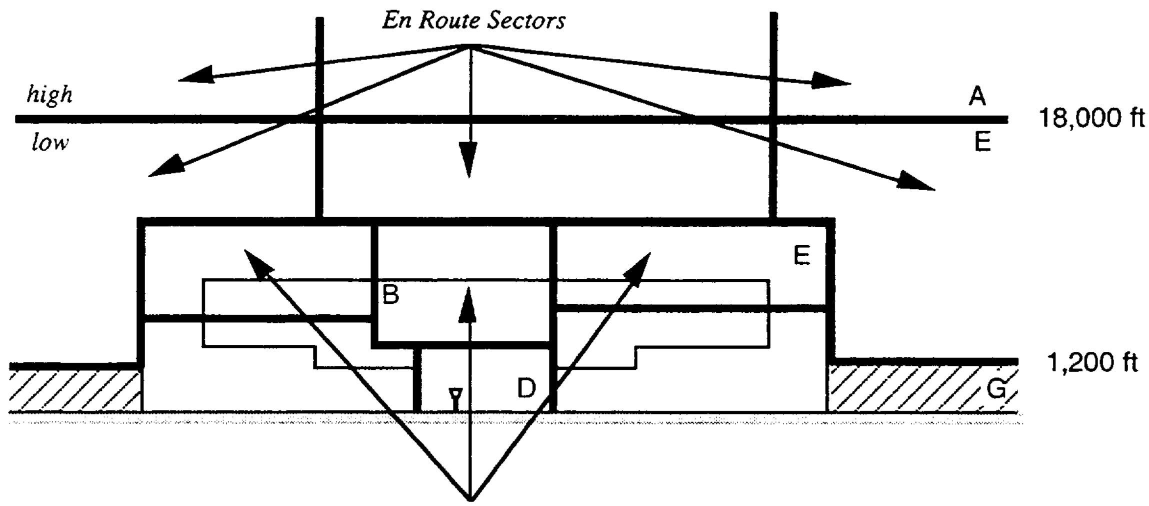

Figure 2.1-2: Terminal Area and En Route Airspace

Controllers authorize flights to proceed with an action by issuing *flight clearances*. Clearances are issued for segments of a flight's itinerary. By controlling traffic at intermediate stages, controllers may tactically react to unforeseen weather and demand conditions.

The fundamental service offered by FAA personnel is safe separation distances among all aircraft within the system. *Separation* is considered to be the essence of ATM, because aircraft separation ensures safety. Separation intervals also determine throughput at airspace sectors and airports. Traffic management specialists continuously study tradeoff methods of increasing throughput without compromising safe aircraft operations. Aircraft separation is divided into three types: longitudinal, vertical, and lateral. Of the three, longitudinal separation has the greatest effect on throughput. Maintaining safe separation distances is complicated by the fact that aircraft travel at different speeds, altitudes, and may be on intersecting or merging routes.

Fundamental to a controller's ability to provide aircraft with separation is the controller's knowledge of each aircraft's location in a sector. Since different monitoring techniques, e.g., radars, transponders, or direct communication with pilots, yield position data with different tracking resolution, FAA separation standards are largely dependent upon the manner in which an aircraft is being controlled.

There are basic operational rules for maintaining required separation minima, along with many exceptions and variations that reflect specific traffic situations. In the terminal area, however, *visual separation* is an alternative means of separating aircraft. Visual rules apply if a pilot can maintain continuous visual contact with the preceding aircraft. Under visual rules, the separation interval in the terminal area may be decreased to as low as 1/2 mile (as opposed to 3 miles under IFR) depending on the type of aircraft involved. If weather permits the use of VFR approaches, an airport's arrival rate should increase, as arriving flights land with reduced separation distances. At the same time, controller workload is reduced, since pilots assume the responsibility of maintaining separation.

Finally, the aircraft ground movements require extensive coordination among Ground controllers (responsible for taxi and ground movement), Local (runway) controllers and pilots. During heavy traffic periods, especially when accompanied by poor weather conditions, ground movements often become an airport capacity constraint. Operational safety concerns also become an issue under these circumstances.

2.2 FREE FLIGHT

The current ATM system is a result of an evolutionary growth of a system over 40 years. While providing high measures of safety, the complexity of the current system does not allow it to easily accommodate changing circumstances in a higher density environment. Direct and immediate gains to the airlines will result from factors including: shorter, more direct routes (no longer limited by the VOR-routes), more efficient arrival and departure procedures, reduced ground and en route delays, operation of aircraft at most efficient speeds, altitudes and in favorable winds, higher airport capacity in all weather conditions and better management of weather related diverts.

A new class of Air Traffic Management systems is gaining acceptance among aviation industry and regulatory organizations. Its fundamental concept is free flight, whose en route embodiment is a safe and efficient operating capability under IFR in which the operators have the freedom to select their path and speed in real time, i.e., User Preferred Routing (UPR). Air traffic restrictions are only imposed to ensure separation, to preclude exceeding airport capacity, to prevent unauthorized flight through special use airspace and to ensure safety of flight. Restrictions are limited in extent and duration to correct the identified problem. In essence, free flight should eventually provide aviation users the flexibility of VFR operation while preserving the traditional safety of the IFR flight.

The key to a much improved efficiency is the need to reduce the rate of conflicts that would require intervention. The FAA believes that the projected traffic growth under the current ATM, coupled with the absence of new automation, could lead to an excessive conflict rate and an unmanageable controller workload. The proponents of the new concept argue that the current funneling of large traffic volumes over discrete geographic points, along with the limitations in radar surveillance, artificially increase the number of conflicts. In free flight the GPS-based aircraft position and velocity determination and their timely communication to the controller are expected to allow a reduction in separation standards, which in turn, is expected to drastically reduce the conflict rate.

In contrast to the current operational paradigm, in the free flight-based system the flight plan "contract" will not be needed. A shift will become possible from a strategic (flight path based) separation to one of tactical (local, near term) separation. Optimum, dynamic flight paths would thus become possible. While the (on-board) Flight Management System-generated flight path will be communicated to the air traffic service provider, it will be used for flow planning and not for insuring separation. The future Automatic Dependent Surveillance (ADS) system will provide the air traffic service provider with accurate position and short term intent information. Advanced automation is essential, to the timely identification of potential conflicts and to the generation of appropriate advisories and resolution instructions.

It is obvious that the equipment required for full free flight operation will not become available to all users of the airspace at the same time. It is crucial to provide an evolutionary path, as a function of technology availability and affordability, procedural changes, aviation community requirements and increase in airspace system capacity. The intent is to make the acquisition of free flight equipage benefits-driven rather than mandated.

The free flight enabling technologies are GPS-based navigation and the faster and more reliable data and voice communications over both line-of-sight (LOS), (i.e., VHF and UHF) and beyond LOS (BLOS) (i.e., HF and especially satellite communications) media. GPS-based navigation, (with appropriate augmentation where needed), will provide much more accurate aircraft position and velocity information, reducing the need for large protective "bubbles" around the aircraft that are required to accommodate large uncertainties in aircraft position. Eventually, it is expected that the accuracy will improve to the point that landing operations in all types of weather will be possible without additional expensive aids such as the currently used instrument landing system (ILS), Automatic Direction Finder (ADF) or VHF Omnidirectional Radio (VOR). In particular, differential GPS, using surveyed locations, appears to be heading towards the required level of precision and dependability.

The gradual shift to digital communications, for both surveillance (ADS) and controller-pilot information exchanges (first the two-way data link (TWDL), to be succeeded by the full controller-to-pilot data link communications (CPDLC)), over the high integrity Aeronautical Telecommunication Network (ATN) is also key to a successful implementation of a free flight concept.

Last, but clearly not a trivial aspect, the ground operations also hold the potential for significant economic improvement. Today's airport surface operations are impaired by limited visual and radar coverage. One advanced concept is to implement a wider use of the digitally transmitted (initially over the Aircraft Communications Addressing and Reporting System - ACARS, later over the ATN) pre-departure clearances (PDC) that would reduce frequency congestion and greatly enhance the accuracy of delivery. The availability of ADS (with suitable augmentation) would allow the controller to "see" traffic in places previously out of coverage. Taxi and takeoff delays would be reduced, increasing the terminal capacity in all weather and visibility conditions to near visual rates.

2.3 Air Traffic Flow Management

Air traffic flow management (ATFM) refers to the management and control of aircraft operating through airports and airspace sectors in a manner that achieves safe, orderly and efficient movement of traffic. ATFM is implemented through a hierarchical command and control system as illustrated in Figure 2.3-1 below. In that figure, the term *Airfield* refers to the physical airport and surrounding ground facilities. The *Local Area* refers to the airspace surrounding an airport, extending out to 5 nautical miles and 2,500 feet in altitude, i.e., Class D airspace.

From the perspective of the FAA, air travelers and the airlines place demands on air traffic resources, such as landing slots at an airport or access to a specific air route. In performing ATFM functions, FAA personnel seek to interfere minimally with the plans and intentions of the airlines. However, in a constrained resource scenario, e.g., too few landing slots available to meet the demand at a destination airport, traffic management specialists must either better manage the resource, thereby increasing its effective capacity, or manage the flow of traffic by adjusting flight itineraries.

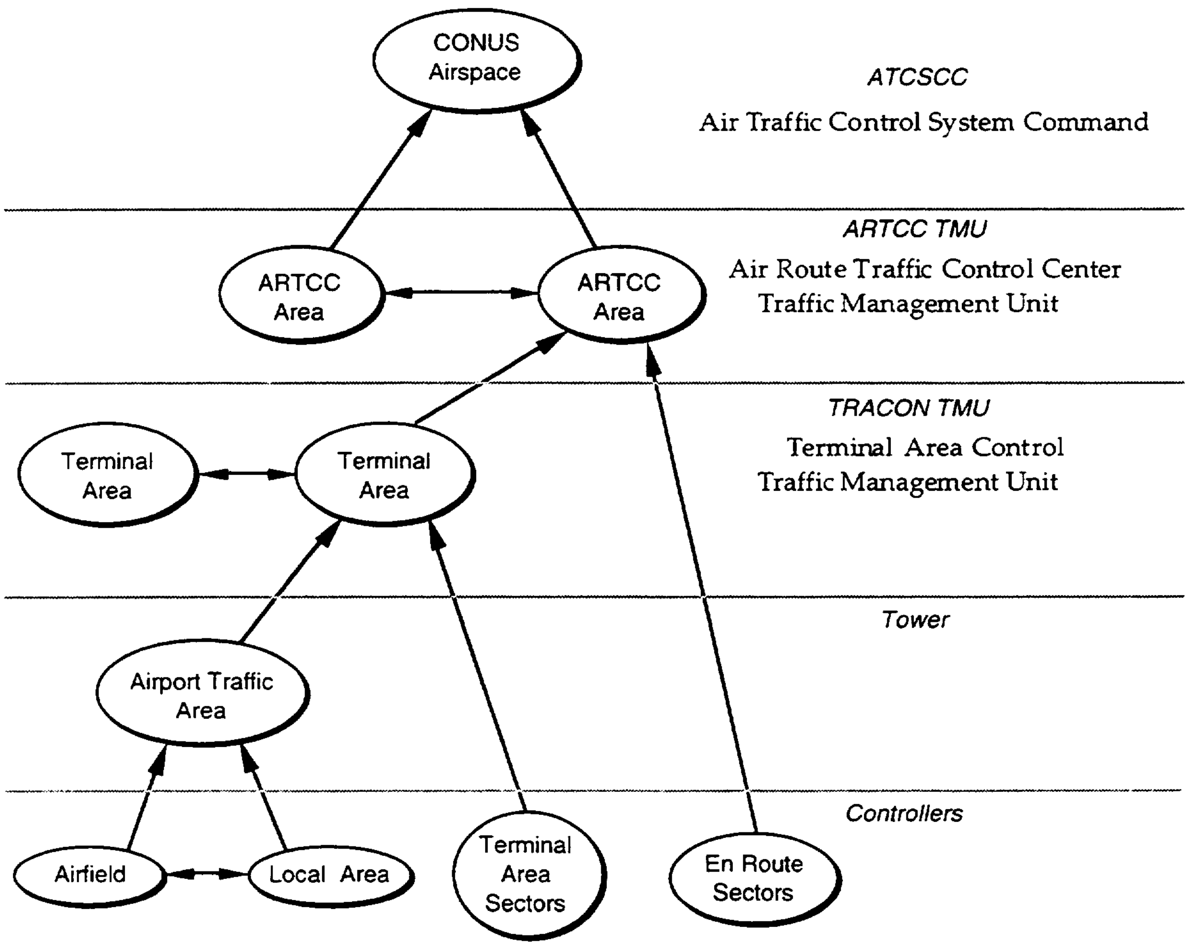

Figure 2.3-1: Airspace and ATFM Hierarchy

2.4 TODAY'S HIERARCHICAL COMMAND AND CONTROL ATFM SYSTEM

The Air Traffic Control System Command Center (ATCSCC), located at Dulles International Airport near Washington D.C., oversees and coordinates ATFM in the Continental United States (CONUS). ARTCCs (Air Route Traffic Control Centers, or simply "Centers") manage en route airspace sectors to ensure safe operations and to reduce potential overcrowding within the capacity-limited terminal area airspace around airports. TRACONs (Terminal Area Control Facilities) are responsible for the management of terminal area airspace by directing departing and arriving flights through the heavy traffic sectors surrounding major airports.

Congestion occurs in the air traffic system when demand for a resource exceeds its capacity. Three air traffic system resources experience congestion: airports, airspace sectors, and fixes. Congestion can cause: high controller workloads, excessive flight delays, safety concerns, and costly, unnecessary fuel consumption. Both Centers and TRACONs employ Traffic Management Units (TMUs) staffed with traffic management specialists who strive to minimize congestion through plans for strategic resource allocation and tactical ATFM directives.

In general, attempts are made to handle flow management problems at the lowest possible level of the ATFM command and control hierarchy. If a flow problem cannot be isolated to a specific region and handled by the regional control personnel or if a traffic problem within a region will ripple into other resources, it will be elevated to a higher level. As problems are elevated, the facility or controller at the lower level(s) remains in the loop of the ATFM decision process. This is possible through extensive communication and data interchange among TMUs, TRACONs, and ATCSCC.

2.4.1 Tower and TRACON

Once in the terminal area, an arriving flight normally will touch down within at most 20 minutes, and hence there are very few opportunities for long-lead forecasting and strategic flow management initiatives. Tower and TRACON efforts are focused on maximizing the use of the terminal area airspace and runways. TRACON TMUs explore alternative methods of prioritizing flights arriving into the terminal area and designing patterns that promote efficient traffic flow.

2.4.2 ARTCC

The principal ATFM concern of Center TMUs is the traffic flow into and out of terminal areas. If at a given airport, arrival capacity is deemed sufficient to handle the projected demand, there is no anticipated congestion. However, if an airport's capacity is forecast to be constrained due, for instance, to weather conditions, the TMU's objective is to adjust arriving traffic so that flights enter the airport terminal area at a rate that matches the expected capacity. Similar ATFM strategies can be applied to fix and sector congestion.

The following options are available to Center TMUs in adjusting flow to deal with constrained capacity.

Miles-in-Trail

Miles-in-Trail (MIT) refers to a minimum longitudinal distance between two aircraft. TMU personnel may restrict aircraft to MIT in order to increase the inter-arrival times of aircraft, thus reducing demand. When applied to an airport experiencing moderately constrained, yet stable capacity, MIT can be very effective. However, if an airport were suddenly to experience an increase in capacity while under MIT restrictions, controllers might not be able to take advantage of the available capacity, due to the aircraft spacing intervals. Moreover, the MIT arrival queue may eventually extend geographically and cause disruptions and delays at departing airports. Indeed, aircraft may be held at a departing airport to meet the MIT restrictions of a destination airport hundreds of miles away.

Minutes-in-Trail

Minutes-in-Trail refers to the longitudinal *temporal* separation among aircraft destined for a constrained resource. Minutes-in-Trail is the most common method of implementing longitudinal separation. It is a safe method of controlling a queue with various types of aircraft traveling at different speeds. The benefits and drawbacks of Minutes-in-Trail are similar to Miles-in-Trail.

Rerouting

Rerouting is a method of redirecting en route traffic onto alternate en route airways to avoid severe weather conditions or overloaded en route sectors. Rerouting requires aircraft to deviate from preferred routes between airports or fixes. TMUs attempt to determine in advance the full impact of rerouting as well as the time when the system is expected to return to normal. TMUs working with weather specialists must determine which routes will be available for use during severe weather conditions, accounting for meteorological changes. Severe en route weather problems affecting multiple Center areas are usually elevated to ATCSCC.

Arrival Metering

Arrival metering is the process of regulating the flow of traffic into a terminal area. The process involves clearing flights over arrival fixes at a rate that matches the capacity of the destination airport. As flights reach outer fixes, controllers calculate cross times for arrival fixes that will ensure desired rates of flow. The calculations may be performed manually or with the Arrival Spacing Program (ASP) available at selected Centers. Controllers delay flights as necessary to meet the desired rate, using holding patterns, or other speed-control techniques. From an en route controller's perspective, it is possible to informally initiate limited airborne holding on flights entering a terminal area, if delays can be maintained below 15 minutes. Finally, arrival spacing provides a TMU with an "inventory" of aircraft in airborne holds to fill the available landing slots that may become available.

Controlled Departure Times (Local)

Controlled Departure Times (CDTs) are assigned at the Center-level using *local* Ground Delay Programs (LGDPs) which adjust flight itineraries prior to departure. In order to control demand at a constrained airport, a Center TMU contacts the airport towers located in the Center area, and instructs tower controllers to delay the departure times of specific flights destined for the constrained capacity airport.

The major benefit of ground delay programs is that delays are transferred from the air to the ground, avoiding potential airborne congestion and reducing en route controller workload. Assigning CDTs is the safest ATFM strategy, although it requires continuous monitoring and adjustment to accommodate the actual capacity or demand variations from predicted levels.

Departure Spacing

Departure spacing involves separating departing aircraft at fixed intervals, usually in response to congestion in the terminal area or immediate en route airspace. The FAA has specific guidelines for departure spacing based on weather conditions and runway layouts.

Traffic Stops

A Traffic Stop holds on the ground indefinitely all flights bound for a problem airport, and hence constitutes a more severe form of CDTs. Traffic stops are initiated when critical events occur at a destination airport. For example, a severe thunder squall with high winds may reduce airport capacity to zero.

CDTs and Traffic Stops are thus the two most severe options available to Center TMUs when managing a constrained resource. The major limiting factor for CDTs and Traffic Stops is that Centers only control the departure times (and hence arrival times at destination airports) of flights originating within their respective boundaries. If CDTs or Traffic Stops are insufficient in relieving congestion, a Center will elevate the problem to the ATCSCC level.

Combinations of two or more of the above options are also possible. The process of selecting the appropriate strategy depends on the characteristics of the problem, i.e., long-term or temporary problem, severe or mild. Two other influencing factors are the on-duty staff level at the TMU and controller experience levels.

2.4.3 ATCSCC

The Air Traffic Control System Command Center (ATCSCC), located near Washington D.C. at Dulles International Airport, oversees all air traffic in the United States. Staffed with controllers and personnel from various sites across the country, the ATCSCC coordinates national flow management initiatives and general air traffic activities. ATCSCC controllers rarely interact directly with individual flights, but rather exert influence on the aggregate flow of aircraft in the system. National flow management actions are initiated at the ATCSCC and transmitted to the appropriate Centers, TRACONs or Towers.

There are five operating divisions within the ATCSCC: Eastern Complex, Central Complex, Western Complex, Severe Weather/National Route Management, and Special Traffic Management.

The Eastern, Central, and Western complexes monitor and assist flow management in the eastern, central, and western United States. Controllers monitor capacity and demand levels of air traffic resources, coordinate flow initiatives among Center TMUs, and, if necessary, administer national FM procedures, i.e., in-trail restrictions, reroutes, ground delay programs and traffic stops.

The Severe Weather/National Route Management division manages severe national weather problems and other large-scale crisis scenarios that disrupt the flow of traffic, and necessitate adjustments in demand and resource utilization. Traffic management specialists continuously research and develop alternative routes for recurring weather problems.

The National Route Management division coordinates and reviews requests by aircraft operators to deviate from preferred routes between airport pairs. Requests are typically received from airlines seeking routes at altitudes with more favorable prevailing wind speeds. Proposed flight trajectories are reviewed for adverse effects on sector or airport traffic levels and are transmitted to Centers managing the sectors affected by the changed itineraries.

Finally, the Special Traffic Management Office handles the coordination and approval of flight plans coordination and approval at the national level. Much of this function is automated.

2.4.3.1 Ground Delay Programs

The purpose of Ground Delay Programs (GDPs) is to coordinate, at a national level, the departure times of flights destined for airports projected to have insufficient arrival capacity (the "problem airports") in order to maintain operationally acceptable traffic levels and reduce airborne holding in airspace surrounding problem airports. In principle, a national GDP resembles a local (Center level) LGDP. However, the scope of a national GDP extends to all flights in the CONUS destined for problem airports. Efforts are made to distribute ground delays in an equitable fashion amongst the airlines and other traffic.

If a Center implements an LGDP or another inter-Center metering strategy, the action is coordinated with the ATCSCC. The problem is elevated to the ATCSCC, if it is determined that capacity constraints or demand levels at an airport warrant a national GDP. On occasion, it is immediately recognized that a problem will be too large to be managed at the Center level and the ATCSCC becomes involved at the onset. The following addresses the two types of GDPs (select and general), the process of implementing a program and follow-on programs. Figure 2.4-1 outlines the typical GDP implementation cycle.

Select Program

A select GDP identifies specific flights for delays based on a unique condition at a problem airport. Select programs are employed when sufficient lead time exists to design the program input parameters around the problem airport condition. The ground delays assigned by a select GDP combined with the scheduled departure times produce new release times for the selected flights, known as Controlled Departure Times (CDTs).

General Program

In general, a GDP will consider all national flights destined for a problem airport as candidates for ground delays. However, if the ATCSCC controllers desire to limit the range of flights affected by GDP delays, general programs may be tailored to exclude flights from certain regions or airlines. The output of a general GDP groups flights together by 15 minute time intervals and each time interval is assigned a *delay factor* by the ground delay program. Delay factors, when applied to individual flights, produce Expected Departure Clearance Times (EDCT) for each flight. The Managed Arrival Reservoir (MAR) allows for up to 15 minutes of delay to be taken in the air and provides a buffer for uncertainties in the system, lessening the chances of underutilized arrival capacity.

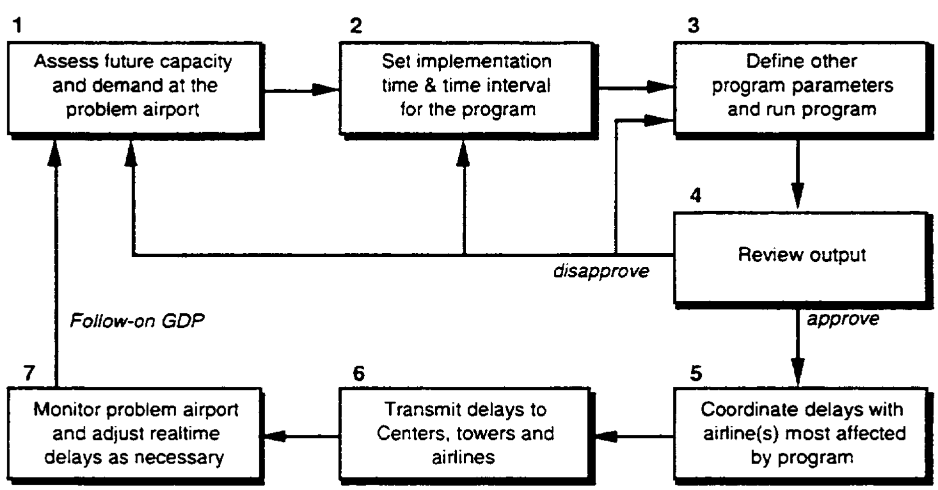

Figure 2.4-1: Ground Delay Program Implementation Cycle

Follow-On Programs

Due to uncertainties in capacity forecasts, traffic management specialists typically limit the temporal horizon of a delay program to four to six hours. However, at times, problems that constrain capacity, e.g., adverse weather, exceed these limits and this requires *follow-on* programs. In order to transition smoothly across programs, follow-on programs are implemented prior to the stop time of the initial program. Follow-on programs often benefit from more accurate capacity forecasts, since the constraining problem at an airport, such as a weather condition, will have materialized during the interval of the initial program. The figure below illustrates the time sequence of an initial program (GDP1) and a follow-on program (GDP2).

Follow-on Program Timeline

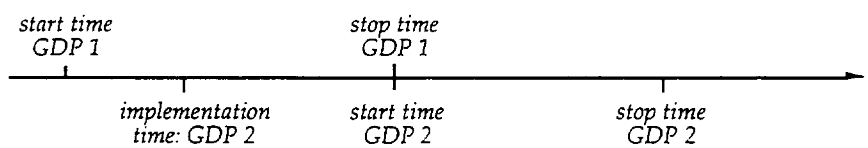

2.4.4 Airlines

Air traffic flow management is next addressed from the perspective of the airlines. Clearly, flight delays are of grave concern to service-oriented and profit-motivated airlines. Although they are aware that the ATCSCC, TRACONs and towers attempt to manage the aggregate flow of air traffic in the network, airlines are nonetheless focused on their individual priorities. Three key airline priorities that are not currently explicitly considered in ATC flow management decisions are the effects of flight delays on:

1. Flight interdependencies.
2. Banks of flights.
3. Flight crew compliance with ATM regulations.

Each is described below.

Flight Interdependencies

A single aircraft typically executes several flight legs throughout the day, creating *flight interdependencies*. Thus, the delay of a particular flight leg may have adverse downstream effects in the form of delays of the same aircraft on subsequent flight legs. The *turn time* of an aircraft is the time required to refuel, execute passenger transfers and re-cater the aircraft. Airlines accommodate turn time by scheduling connecting flights with sufficient layovers, or *ground times*. The amount of slack in the ground time in comparison to the turn time (what might be referred to as *ground slack time*) determines whether or not a small delay on an arriving flight may result in a late departure for the connecting flight. Airlines look downstream and evaluate the ability of an aircraft to absorb, through ground slack time, a delay in the future when reviewing and accepting the Controlled Departure Times transmitted from the ATCSCC.

Banks of Flights

The hub system, adopted by the major carriers in the 1980s, further complicates the effect of flight delays. Under the hub system, an airline schedules many flights to arrive into a hub airport within a short interval of time. Such a group of flights is referred to as an *arrival bank*. Following a short time period to allow for passenger connections and aircraft preparation, the airline schedules a *departure bank* of flights. A major weakness in the hub system is its sensitivity to flight delays. Specifically, ground delays or airborne holds imposed on one or two flights in an arrival bank, may result in airline-imposed delays on multiple flights in the related departure bank. Airline decisions to delay departing flights are based on the percentages of connecting passengers on the late arriving flights and, as mentioned above, potential downstream effects.

Flight Crew Compliance

Delays can also have a flight crew-related effect in that crews must abide by FAA regulations relating to the maximum allowable duty time. Unplanned lengthy delays may cause a crew to reach its maximum allowable duty time prior to completing all scheduled flights. When this occurs, airlines react by adjusting crew schedules where possible or by canceling flights. In addition, crews do not always continue on the same aircraft so that a late arriving aircraft may hold up a different aircraft waiting for its crew.

3. Opportunities for Improvement of the Existing ATFM System

Here we consider two principal classes of problems in ATFM: those related to issues of capacity and its efficient use, and those related to issues of safety. The first class is addressed directly. Safety is a constraint that must be imposed on any solution proposed to address efficiency related problems and, therefore, will be attended to at least implicitly, if not explicitly, in the efforts proposed here. This section serves as an introduction to some of the areas where there are opportunities for improvement to the existing ATFM system.

The capacity of the air transportation system depends on a variety of factors, some that cannot be altered by human action and some that can. The primary factor in determining capacity that cannot be altered by human action is weather. Those factors that can be affected by human action include the physical resources of the system (e.g., number of runways at an airport) as well as the technologies and procedures employed in using these resources (e.g., CTAS). Modifications to existing ATFM technologies and procedures that are focused on increasing capacity and its efficient use—along with the introduction of new technologies required to support those modifications—represent opportunities for potential system-wide improvement in air transportation flow management.

3.1 Increase System Capacity

The most direct approach to improving the existing ATFM system would be to increase the capacity of the air transportation system by building more airports and/or runways. As evidenced by the problems attendant with the recent opening of the Denver International Airport, there are significant financial and political pitfalls in this.

Another strategy for improving capacity would be to increase the effective capacity of the system's resources through changes in technologies and procedures whose objectives are the more efficient use of those resources. For example, the effective capacity of the *terminal area airspace* can be increased through better planning and controller decision support (e.g., CTAS).

The effective *runway capacity* could be increased by employing technology to improve ability to land in bad weather (Cat III or autonomous landing). The *ground traffic* throughput capacity could be increased by improving the ability to plan, manage and monitor ground (taxiway) traffic in bad weather through advanced surveillance technologies, onboard displays and better tower decision-support. The effective capacity of en route sectors could be increased by not limiting traffic to traditional airways, e.g., through the introduction of free flight.

Although capacity-related problems represent the area for greatest potential improvement to the ATFM system, there remain a number of issues that must be addressed in order to realize the maximum possible improvements in the ATFM system. Three of these are:

(a) a proper understanding of how uncertainty impacts planning and execution of ATFM directives

(b) resource interdependencies that must be accounted for properly in order to insure that correcting one problem doesn't exacerbate another and

(c) system-wide metrics that must be employed in determining whether a change to the ATFM system provides an overall improvement in system performance.

Each of these is discussed in more detail in the following.

3.2 IMPROVING KNOWLEDGE OF UNCERTAINTY

There is significant uncertainty in the ability to predict both the capacity of and demand for system resources.

3.2.1. Uncertainty in Capacity

As we have seen, predicting the capacities of and the demands for system resources forms the basis for the ATFM planning process. Furthermore, the accuracy and lead times of those predictions are critical to successful outcomes in executing ATFM plans. Resource capacity reductions are primarily due to the influence of weather on air traffic operations. For example, low visibility in the terminal area requires increased aircraft spacing intervals as well as degraded taxi and ramp traffic conditions, both of which decrease traffic arrival capacity of an airport. Difficulties in predicting a resource's capacity even a few hours in advance can result from the uncertainty in forecasting the weather conditions that directly impact those capacities.

In order to minimize the uncertainty in these predictions, traffic management specialists work closely with weather personnel when projecting capacity forecasts prior to developing and implementing plans for ATFM initiatives. If a prediction underestimates a resource's capacity, actions will have been initiated to avoid exceeding the inaccurately predicted, low level of capacity and the resource will possibly be "starved" of aircraft over the horizon of that prediction. Conversely, if a resource's capacity is overestimated, the resource may become constrained requiring real-time ATFM interventions which often result in undesirable airborne delays.

In many cases, the lead time of a forecast is as important as its accuracy. In particular, effective ATFM requires that traffic managers have advance warning of potential resource capacity reductions that is sufficient to allow for planning and initiating actions in response to those reductions. Indeed, perfect knowledge of the status of the system even one hour in advance may provide insufficient lead time to implement an ATFM action required to relieve congestion for some airports or sectors.

3.2.2. Uncertainty in Demand

In addition to resource capacities (i.e., supply), the demand on airports, airspace sectors and fixes is also uncertain. Unannounced VFR traffic is one contributing factor. The first indication of a VFR flight's intention to enter a terminal area, may be an in-flight contact by the pilot requesting clearance. General aviation traffic under IFR may also arrive unannounced into a congested area (a "pop-up") if the flight plan were not introduced into the system with sufficient lead time.

The durations of the various components of a flight's planned trajectory are themselves subject to uncertainty which, in turn, affects the estimates of the future demand that the flight will place on resources. For example, departure congestion may result in a delayed flight departure, shifting the schedule of that flight for its entire trajectory so that the aircraft reaches fixes, sectors and the destination airport each at times later than the scheduled times. En route flight transit times can be uncertain as well, being influenced by the prevailing jet stream conditions. A high speed jet stream may cause an early arrival for flights traveling west to east, or a late arrival for flights traveling east to west. In both scenarios, demand projections become uncertain.

Further uncertainty in demand is introduced by real-time airline practices. Airlines, notwithstanding published schedules, react to the daily air traffic environment (weather, delays, etc.) and the air travel market. In particular, airlines may adjust flight times or cancel flights based on excessive delays, low passenger demand, crew availability, equipment problems and/or scheduling constraints.

There are several opportunities for improving the existing ATFM systems by better approaches to dealing with uncertainty. The first is to decrease uncertainty by more accurate and more timely (i.e., better lead time) predictions of weather. The second is to employ approaches to planning for the utilization of future expected capacity that *explicitly* account for the nature of the uncertainty, i.e., stochastic optimization approaches that are capable of taking advantage of knowledge of the fact that the future is uncertain and of the characterization of that uncertainty. These approaches are designed to do the best job at managing the tradeoff between being too conservative and underutilizing (on average) resources and being too optimistic and causing costly, less safe airborne delays due to rerouting or airborne holding patterns.

3.3 Modeling Resource Interdependencies

In addition to the considerations discussed in the preceding, to be effective from a system-wide perspective all solutions to capacity related problems must explicitly model and account for what has been referred to earlier as resource interdependencies.

Even with perfect predictions well in advance, congestion may be unavoidable due to resource interdependencies. If one resource is identified as facing potential congestion, a traffic management specialist must identify an alternative resource, or resources, to either hold or absorb the excess demand. If alternative capacity does not readily exist, congestion may ripple through the system. For example, one interdependency is the relationship between the arrival capacity and departure capacity at an airport. For a given runway configuration, there is a tradeoff between the number of arrivals that an airport can accommodate versus the number of departures, referred to as the arrival / departure frontier (see Figure 4.3.3-1). If departure congestion at an airport becomes excessive, the tower may choose to favor departures over arrivals to "clear out" a ground departure queue, thereby delaying planned arrivals due to reduced arrival capacity.

Efficiency-related problems are those that result from difficulties in managing a system that has *significant* temporal and spatial excess demand (e.g., reduced capacity due to weather). Recall that the principal air traffic system resources with limited capacity are en route sector capacity, terminal airspace capacity, landing and takeoff capacity (and the associated tradeoff between the two) and the ground traffic movement or throughput capacity. Gate and ramp traffic capacity are not addressed since control of those resources (in the US, at least) is the province of the airlines.

3.4 SYSTEM-WIDE METRICS OF EFFICIENT OPERATION

In order to understand the effect of changes in procedures and technology, there must be system-wide metrics for measuring the efficiency of operations that are acceptable to all of the those participating in the use of the system: the FAA, the airlines and the passengers. Indeed, efficient use of the capacity may mean different things to the different players in the system. To the FAA, efficiency might be measured in terms of the proportion of the system capacity that is actually used to meet demand. To the airlines, efficiency might be measured in terms of the profitability of operations. To the passengers, efficiency might be measured in terms of getting to the destination quickly and reliably at a reasonable cost.

Any modifications to the procedures of the current ATFM system that address efficiency problems must be designed and evaluated to ensure that, in the process of improving efficiency of one aspect of air transportation system operations, safety is not jeopardized and broader system-wide capacity-related problems are not exacerbated. Indeed, ATFM decisions that ameliorate capacity-related problems for a given air transportation system resource may have a negative impact on the efficiency of the operations of one or more individual airlines. Thus, any modifications to existing ATFM capabilities and practices must take into account the different players' metrics of efficiency, requiring the development of system-wide metrics that address the concerns of all players. A fair resolution of this issue is at the heart of the challenge imposed in designing, developing and evaluating concepts for changes to the air transportation system that yield truly system-wide improvements in operations.

4. ALTERNATIVE ATFM CONCEPTS

As long as there are times when the capacity of one or more air traffic system resources falls considerably below the scheduled demand such as during severe weather conditions, a coordinated approach will have to be employed to allocate fairly scarce resources during the period of excess demand. Indeed, it is highly probable that there will continue to be significant periods when demand for the limited departure and arrival capacities at the busiest airports in the current air traffic system exceeds the supply.

In a *centralized* ATFM system during periods when demand exceeds capacity, the ATFM system's operator makes largely unilateral decisions regarding the assignment of delays to every aircraft, the modifications of their routes, etc. and monitors closely each aircraft's compliance. A *moderately decentralized* concept would be similar in many ways to today's ATFM system but with increased FAA / airlines cooperation and coordination in ATFM decision-making. In a *fully decentralized* ATFM system, each aircraft and aircraft operator would be given accurate and timely information about existing and projected demand and capacity for each ATFM system resource, allowing each aircraft operator to determine independently its own preferred strategy with regard to its own set of flights, and to plan, execute and monitor its own detailed commands, with the ATFM system operator kept apprised of those plans and intervening only when needed for safety.

The airlines have generally indicated a preference for the decentralized end of the spectrum. Indeed, under more decentralized schemes, airlines, in theory, would have more freedom to optimize their individual operations, with the potential for providing passengers with more efficient and more reliable travel, as measured by shorter flight times and more reliable schedules and connections. The free flight concept is an expression of this theory (RTCA, 1994). The Collaborative Decision Making (CDM - formerly called FADE) program (Wambsganss, 1995), already under way, is a first step in the attempt to increase cooperation and collaborative decision-making between FAA and airlines in the existing ATFM system.

4.1 Evolution of Alternative ATFM System Concepts

The ATFM system is an extremely complex, large-scale system that can be decomposed into three highly coupled physical segments — en route, near terminal area and ground operations (see Figure 4.1-1). The decisions made for planning and controlling the traffic management activities within each of these segments impact those within the other two, and together they impact the overall flow of traffic through the air transportation system network. Ultimately, the objective of any modification made to the ATFM system is to increase the effective capacity of the overall system in ways that benefit all participants (FAA, airlines, general aviation and passengers) while sustaining or improving the level of safety afforded by the system. Because the decisions and activities within one segment impact those in the others, some form of coordination of air traffic management across segments will be required both to increase opportunities for synergism across segments that lead to improving the overall system-wide performance and to insure that solutions to air traffic management problems in one segment do not have a significant negative impact on other coupled segments. The metrics or figures of merit that will be applied in evaluating alternative ATFM concepts are discussed in Section 5.

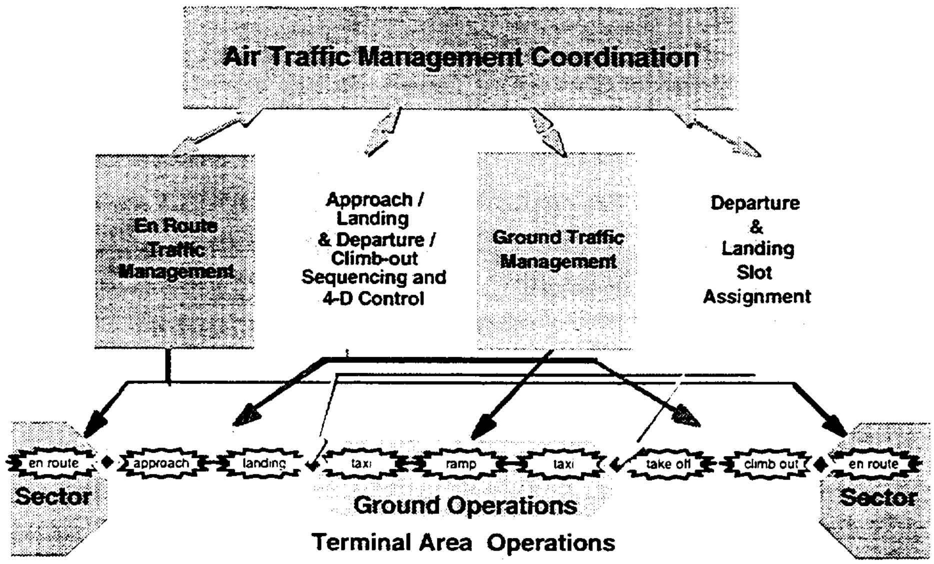

Figure 4.1-1: Coordination of Traffic Flow Management across Segments

Figure 4.1-1 represents a functional decomposition of decision-making in ATFM, but does not reflect an allocation of those decision-making functions to the FAA, airlines or any other potential participant in the system.

In this report, the term *decentralization* refers to both the decomposition of the decision-making functions as illustrated in the middle tier of Figure 4.1-1, as well as the shared allocation of those decision-making functions between the FAA and airlines. In the latter case, decentralization connotes that the airlines play a more significant role than heretofore, resulting in reducing the current level of centralization of decision-making.

The evolution from the policies and procedures under which the current ATFM system operates to those of future, more decentralized systems must occur along a migration path of feasible, cost-effective changes. Each change along such a path should provide improvements in system performance. Specifically, new equipment and procedural changes must be phased into the system in such a way that there are benefits to the overall system as measured by its benefits to the individual participants: the FAA, the airlines and the passengers.

Feasibility of implementation is a criterion that must be applied in defining the sequence of system modifications. In particular, the introduction of new technologies cannot occur "overnight" so that the system must operate under each modification with a mixture of equipped aircraft, some of which will have new equipment and some of which will not.

4.2 Characterization of Alternative ATFM Concepts

Table 4.2-1 outlines a spectrum of alternative ATFM system concepts, with each row representing a potential operational alternative. A dynamic, real-time decision-making environment is assumed throughout. The columns in the table decompose the strategies for implementing each concept into:

1. Initial allocation of arrival slots among aircraft operators.
2. Final assignment of arrival slots to individual flights.
3. Assignment (if any) of departure slots to individual flights.
4. En route flight planning and control.
5. Transitional area and terminal area flight planning and airport surface movement control.

The concepts in the table range from "highly centralized" in the top row (where decision-making is centralized within the FAA) to "highly decentralized" in the bottom row (where decision-making is decentralized across the airline operators). Moving down a row in the table represents an evolutionary step in the process of system decentralization and a consequent move in the direction of free flight (RTCA, 1994). Each alternative ATFM operational concept (each row) must be analyzed and evaluated to find the best operating point from the point of view of both system performance and system safety.

The second row, *Partially Centralized*, roughly describes the present state of the system. The third row, *Partially Decentralized I* — highlighted in gray — is the focus of the discussion in Section 5.1. The ATFM concept represented by the third row is a viable evolutionary step from the present system; indeed, most of the technologies to implement this concept exist, and consensus for it on the part of the FAA and airlines is beginning to take shape.

	Allocation of Arrival Slots	Assignment of Arrival Slots to Individual Flights	Assignment of Departure Slots to Individual Flights	En Route Planning and Control	Transition Area, Terminal Area, Ground Movement Planning and Control
Centralized I	1a TFM system operator (FAA) allocates arrival slots to individual flights.	2a TFM system operator assigns slots; each airline may cancel and substitute flights	3a TFM system operator assigns departure slots to individual aircraft.	4a Airlines plan; TFM system operator controls.	
Partially Centralized (Current system)		2b Each airline suggests alternative assignment of the slots allocated to it (for its own flights only); TFM system operator approves or rejects.			
Partially Decentralized I	1b TFM system operator allocates sets of arrival slots to individual airlines.	2c Individual airlines allocate their own sets of slots among their own flights.	3b Airlines assign departure slots to individual aircraft; TFM system operator approves or rejects.	4b Airlines plan; TFM system operator specifies a region in which airlines control their own aircraft; TFM system operator controls other regions and monitors globally for feasibility, conflicts.	5a Airlines plan; TFM system operater controls.
Partially Decentralized II		2d Airlines may trade slots among themselves. Each individual airline allocates its own set of slots among its flights.			
Decentralized I	1c TFM system operator informs airlines of the legal safe capacities at potentially congested airports.	2e Airlines may bargain among themselves for the legal safe capacities. They may cancel or delay flights, follow the original schedule, etc. within their "purchased" slots.	3c Airlines assign departure slots to individual aircraft.		5b Airlines plan; they also specify each aircraft's heading directly after departure; TFM system operator can approves or rejects heading; TFM system operator controls everything else.
Decentralized II	1d TFM system operator informs airlines about anticipated availability of capacities at congested airports.	2f Airlines decide what they will do. They may cancel or delay flights, follow the original schedule, etc.		4c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.	5c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.

Table 4.2-1 Spectrum of Alternative ATFM System Concepts

All concepts outlined below the *Partially Decentralized I* row in Table 4.2-1 should be considered *highly speculative* at this time and are listed here only as *a rough indication of the types of potential approaches that may emerge in the future*. The brief descriptions of these concepts in Section 4.3 below are necessarily sketchy and incomplete. Moreover, to our knowledge, no analysis of the *"robustness"* of these concepts with respect to ensuring a fail-safe system operation has been performed to date. Thus, the feasibility of moving beyond a concept similar to *Partially Decentralized I* is an open question.

4.3 Discussion of Table Entries

4.3.1. Column 1: Allocation of Arrival Slots

One of the principal flow management strategies employed in the current system in reacting to reduced capacity at an individual destination airport is that of controlling the departure times of aircraft destined for that airport through the use of Ground Delay Programs. In effect, these controlled departure times implicitly represent a set of arrival slot allocations at the destination airport of interest. The current strategy employed in arrival slot allocation is one wherein the FAA attempts to minimize airborne delays and maximize utilization of available airport take-off and landing capacity with the objective of adhering as closely as possible to published schedules (i.e., OAG). Recently, in order to ensure that airport arrival capacity is not wasted, a program referred to as Managed Arrival Reservoir (MAR) has also been instituted. This program allows up to fifteen minutes of arrivals beyond the forecasted capacity to be airborne (an airborne reserve).

In a more decentralized ATFM system, where free flight-like concepts are employed for en route traffic, the control of the allocation of arrival slots will be the FAA's principal flow control mechanism, and, consequently, is likely to play an even more significant role in air traffic management than it does in the current system.

In a first step toward decentralization (**1b** in the table), the FAA could allocate a set of slots to each airline over predefined intervals without assigning specific flights, giving the airlines more flexibility to assign flights to the allocated slots. This has an advantage from the airlines' perspective in that each individual airline can assign flights to its allocated slots based on its own business objectives.

In the next step, no specific landing slot assignments are made to individual airlines (**1c**). However, to ensure safe operation, there is a cap on the total number of slots for predefined time intervals. This opens the possibility to create a "market" within which the airlines "trade" for slots up to the specified limits. Finally in a fully decentralized scenario (**1d**), the market of available slots is not constrained, but the airlines are kept informed by the FAA of the expected safe limits on arrival capacities at individual airports.

4.3.2. Column 2: Assignment of Arrival Slots to Individual Flights

Under the approach taken today (**2b**), each airline, within limits, suggests alternative assignments of the slots allocated to it (its own flights); the ATFM system operator may approve or reject those alternative assignments. The airlines have the freedom, within established constraints, to cancel flights and substitute other flights. Substitutions are done one at a time; each is subject to approval by the FAA.

2c: Individual airlines freely assign their own sets of slots among their own flights. In this case the FAA performs no assignment of flights to slots. Each airline is allocated slots and is given the freedom to assign any of its own flights to the allocated slots, with the assignments being subject to FAA approval.

2d: Airlines are allocated slots and may trade slots among themselves. Each airline is initially allocated an unassigned set of landing slots at capacitated airports, and the airlines are free to barter among themselves in order to re-allocate those slots. Airlines are free to assign flights to their allocated/bartered slots as they see fit.

2e: Airlines may negotiate among themselves for slots within legal, safe capacities. They may cancel or delay flights, follow the original schedule, etc. within their negotiated slots. The negotiations, for example, could be based on a market or on bartering. In this case, no initial allocation of slots is made by the FAA; rather the FAA sets a limit on legal slots for each airport as a function of time, based on considerations of safety. The total set of negotiated slots, across all airlines, must not exceed the aggregate legal, safe limit established by the FAA. Airlines may launch flights that exceed this limit, but such flights are subject to diversion if they cannot be handled safely at the destination airport.

2f: Airlines decide freely on slot allocations. Airlines may cancel or delay flights, follow their original schedule, etc. In this case, each individual airline makes its own decisions as to how and when to assign flights. The FAA would disseminate information with respect to "safe" capacities and expected aggregate demand based on the most up to date information regarding the assignment decisions of all of the airlines.

4.3.3. Column 3: Departure Slot Assignment

Given an assignment of flights to arrival slots at destination airports, departure slots for those flights from their origination airports can be assigned in a variety of ways. The most straightforward is simply to subtract the nominal flight time from the arrival slot assignment time and assign a departure slot at the originating airport for that time. Alternatively, in situations where the assigned arrival slot for a given flight represents a delay with respect to the scheduled arrival time for that flight, the airlines or the system operator may choose for that flight to depart earlier than the simple difference between the assigned arrival time and nominal flight time (e.g., as in the MAR program) in anticipation that either (a) weather may improve and the increased realized capacity will result in newly opened slots that a "flight already in the air" could take advantage of or (b) a slot may open up at the arrival airport as a result of other flights slotted to arrive being delayed or canceled. Thus, a buffer of aircraft in the air representing demand that slightly exceeds anticipated capacity ensures that an unexpected increase in realized capacity will not go unused. Given the freedom to make this kind of decision, airlines may decide to take some of the scheduled delay on the ground (a ground hold) and some in the air, anticipating that capacity will improve with some non-zero probability.

3b: The airlines assign departure slots to individual aircraft, and the ATFM system operator approves or rejects those assignments. Thus, the airlines would have the opportunity to anticipate potential improvements in arrival airport capacity and would be allowed to choose to "leave early" in order to take advantage of any realized improvements in arrival capacity. To ensure that the number of airborne aircraft is not so large as to cause potential safety problems due to unacceptably high levels of congestion near arrival airports, the chosen departure slots would be subject to approval by the system operator.

3c: The airlines assign departure slots to individual aircraft. Here, the system operator would only be responsible for controlling departures to ensure safety. There is, of course, no guarantee that a flight will be able to depart at the desired time, if the total requested number of departures from the originating airport exceeds the departures capacity of that airport during the period of interest.

For given weather conditions for every airport, the number of departures per unit time interval and the number of arrivals over that same interval are often coupled, as notionally illustrated in Figure 4.3.3-1. Thus, arrival slot allocation and departure slot allocation cannot be performed independently when demand exceeds the available supply.

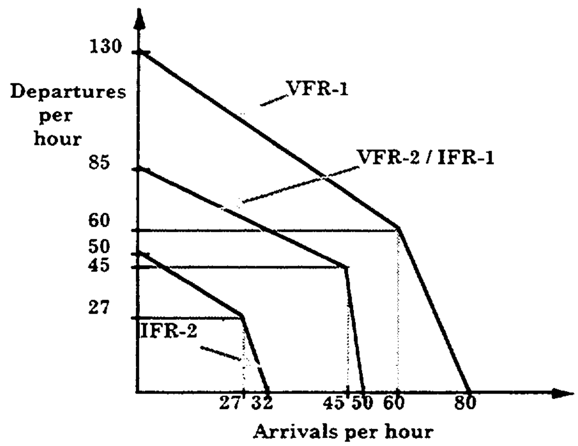

Figure 4.3.3-1: Typical Trade-Off of Arrival vs. Departures

4.3.4. Column 4: Enroute Planning and Control

For the en route segment, control strategies range from assignments of 4-D waypoints by the system operator to free flight.

4a: Airlines (and other airspace users) plan; the ATFM system operator controls. Airlines file a flight plan for each flight and the system operator suggests or mandates modifications to the plan, if deemed necessary for safety purposes, and monitors the flight. Flight plans are expected to largely conform with established airways.

4b: Airlines (and other airspace users) plan; the ATFM system operator specifies regions in which user-preferred trajectories are acceptable, controls flights in other regions and monitors globally for feasibility and conflicts. This is a mix of the approach described in 4a above for certain regions (e.g., in highly congested en route sectors and in near terminal areas) with more freedom for the airlines to fly in "free flight" in other regions. The current National Route Program (NRP) represents the first step in this direction. In all cases, the system operator is required to monitor flight trajectories in order to predict conflicts and to intercede and resolve conflicts when potential conflicts are detected.

4c: Airlines (and other airspace users) plan and control their own flights; the ATFM system operator monitors for feasibility and conflicts. Here the airlines are afforded the flexibility to freely choose en route trajectories, with the system operator responsible for conflict monitoring and, when necessary, resolution.

4.3.5. Column 5: Transition Area, Terminal Area, Ground Movement Planning and Control

In order to ensure safe operations under all of the alternatives discussed here, the system operator must monitor and control traffic in the near terminal area for arrivals and, to a lesser extent, for departures. Under current approaches to ATFM, the system operator implicitly controls the demand on near terminal area airspace through both arrival slot allocation and en route traffic control. With the evolution toward increased freedom on the part of the airlines to choose arrival and departure slots and en route flight plans, there comes the increased potential for substantial congestion in the airspace near airports. To avoid the associated potential for terminal area controller overload and attendant risks to maintaining safe operations, incentives that serve to coordinate ATFM between the en route and terminal areas must be applied by the system operator in order to influence traffic behavior in a manner that reduces this potential.

Although the terminal area traffic may be tightly controlled by the system operator, this does not preclude collaboration by the airlines in creating the plans for the sequencing of that traffic. Currently, arrival sequence planning under even the most advanced systems (e.g., CTAS) does not directly include participation by the airlines. Increased decentralization may provide opportunities for the participation of the airlines in formulating the objectives employed by the system operator in choosing arrival sequences.

5a: Airlines and other airspace users plan their activities within the terminal area and on the ground; the ATFM system operator controls aircraft in this region (except when they are on the ramp) including movement on the taxiways, aircraft departures, climb-outs and descents. Departure headings are limited to predefined fixes.

5b: Airlines (and other airspace users) plan; they also specify each aircraft's heading directly after departure (here we assume that departures are not restricted to flying departure fixes); the ATFM system operator can approve or reject headings. Prior to departure, airlines and other airspace users inform the ATFM system operator of their preferred departure heading. To ensure that the configuration of departing airborne aircraft does not cause potential safety problems, the preferred departure headings would be subject to approval by the system operator.

5c: Airlines and other airspace users have the freedom to operate within the "rules of the road," similar to the way automobiles operate. The ATFM system operator monitors for feasibility and conflicts and has the capability to impose control, just as police authorities do for automobile traffic.

In addition to arrival and departure sequencing and departure heading selection, another component of terminal area ATFM is ground traffic management. In ground operations, there can be a strong coupling between the assignment of departure slots and the control of the ground movement of aircraft. Specifically, in poor weather with limited visibility, arrival and departure rates (airport capacities) can be limited by the ability to move ground aircraft traffic through taxiways and ramps in a timely fashion. Ground traffic planning, management and control can benefit from timely and accurate traffic location information and digital data communications. Thus, improvements in ground traffic surveillance and planning can effectively improve the capacity of some airports in poor weather. Again, as in the case of arrival sequencing, decentralized approaches to traffic flow management must afford the airlines the opportunity to influence ground traffic plans to the extent that those plans control the sequence of departures at an airport.

5. EVALUATION OF NEW ATFM SYSTEM CONCEPTS

Since the ATFM system consists of several highly coupled segments, as described in Section 2.1, it would be misleading to evaluate the impact of a modification to ATFM operating procedures for those segments without accounting for its impact on the rest of the system. In this section an ATFM simulation testbed that embodies a variety of system-level modeling and analysis tools is described. This testbed has been developed to evaluate the system-wide impacts of candidate modifications to the existing ATFM environment.

Since each of the three principal "stakeholders" in the system—the FAA, the airlines and the traveling public—may have different sets of priorities and objectives, each ATFM system concept should be evaluated using a variety of metrics that reflect the "utility functions" of each of the stakeholders.

> The use of a simulation testbed that contains a complete system description, appropriate evaluation metrics, models of each of the system entities and an array of analytical capabilities, will ensure that proper system-wide evaluations are performed.

5.1 A VIABLE MEDIUM-TERM PARTIALLY-DECENTRALIZED SCENARIO

Our research has concentrated on the development of a set of tools and experiments that make it possible to evaluate the "Partially Decentralized I" concept highlighted in Table 4.2.1. As noted previously, this concept presents a viable evolutionary step from the present system. Indeed, some of its aspects will be implemented within the next 5 years under the Free Flight Action Plan recently announced by the FAA (1996) in response to the work of RTCA Task Force on Free Flight, while several other aspects have been investigated, at least in a preliminary manner, in a number of recent studies (Milner (1995), Wambsganss (1995), DeArmon and Lacher (1996)).

The principal characteristics of the concept can be summarized as follows:

1. For slot-allocation purposes, the busy hours of the day are subdivided into intervals that can accommodate several arrivals (e.g., intervals of 10 or 15 or 20 minutes). This increases flexibility for both the system operator and the airlines with regard to dynamic arrival scheduling, while providing protection from excessive arrival clustering (DeArmon and Lacher (1996)). (It should be noted that none of the operational characteristics 3-8 listed below depends critically on this point and that the proposed concept is also compatible with a slot allocation system that would allocate slots on a one-flight basis, i.e., by subdividing the time axis into intervals of the order of 1 minute.)

2. Allocation of arrival slots at congested airports is performed on a dynamic basis, according to predicted airport capacity over the next few hours. Whenever arrival capacity at one or more airports is predicted to be scarce, available slots for arrivals at these airports are allocated among the airlines on a First-Scheduled, First-Served (FSFS) basis to ensure fairness. For example, suppose that 15-minute slot-allocation intervals are in use and that airline XYZ originally had 6 arrivals scheduled to arrive at a particular airport between 10:00 and 10:14. On a day when capacity is low, XYZ might then receive 4 slots on a FSFS-basis, with the other two slots moved to the 10:15 -10:29 interval. Note that, while the number of slots that XYZ will receive in the interval is specified by the ATFM algorithm, the identity of which of XYZ's flights will occupy these slots is not. This is a fundamental aspect of this partially decentralized concept and applies irrespective of whether the intervals into which slots are allocated is 1-, 10- or 15-minutes long.

3. Each airline (or, more generally, each aircraft operator) is free to utilize its slots in each interval in the way it deems best. Thus, each airline may schedule any one of its flights into any one of its arrival slots. Each airline must also keep the ATFM operator informed as to which flight has been assigned to each slot and, most important, as to what slots, if any, will be left unused, due to flight cancellations. Any slots left unused by a particular airline will be awarded by the ATFM operator to other airlines on a FSFS basis.

4. The ATFM operator then estimates a "controlled time of arrival" (CTA) for each flight, (taking into account each airlines announced preferred sequencing of its own flights) and sends to each airline the list of that airline's CTAs. The point of "arrival" is not necessarily the runway; in fact, in the presence of congestion, this point will usually be the boundary between en route airspace and the transitional airspace into the terminal area of each congested airport.

5. Little or no use is made of departure slot assignments (known currently as EDCT, Expected Departure Clearance Times). Thus, each aircraft operator is responsible for determining the time of take-off which is most appropriate for meeting the assigned CTA of each one of its flights. This means that the aircraft operators will also decide how to best allocate any predicted delay resulting from the assigned CTAs between delay taken on the ground and delay taken while airborne. In other words, the airlines will determine the size of their own "Managed Arrival Reservoirs."

6. Free flight (user-preferred routing) is permitted in large portions of en route airspace and is utilized to travel from the airport of origin to the "arrival point" for which the CTA has been specified.

7. Air traffic management in the transitional area to the airport of arrival, in the airport's terminal area and on the airport's surface is supported by advanced decision support and automation aids such as CTAS and SMA.

8. The ATFM operator checks for compliance on the part of aircraft operators with slot allocations and with CTAs. The ATFM operator also monitors continuously traffic operations to ensure safety.

Numerous additional details can be specified with regard to the partially decentralized concept described by 1-8, but are superfluous for the purposes of this report. It should also be noted that many other plausible variations on the above themes merit investigation. The important point, however, is that a consensus is beginning to take shape on the part of the FAA and the airlines about the desirability of moving toward a system that complies with the general framework outlined here.

5.2 EVALUATION METRICS

A set of key metrics related to congestion, delay costs, schedule reliability and utilization of aircraft and other resources has been identified as appropriate for the evaluation of alternative ATFM systems. Each of these metrics quantifies an aspect of performance which is of particular interest to one or more of the three principal stakeholders identified in the previous sections. Fine granularity metrics are listed in Table 5.2.1; aggregate metrics can be derived from those:

CONGESTION	(1) Demand to capacity ratios at each airport and each defined airspace region by time of day; (2) Number of aircraft held/delayed in the air on arrival for each airport by time of day; (3) Number of aircraft held/delayed on the ground on departure due to ATFM intervention; (4) Spatial density of airborne aircraft by defined airspace region over time.*
DELAYS AND DELAY COSTS	(5) Minutes of delay incurred per operation (e.g., arrival , departure, taxi-in, taxi-out, transit of a portion of the flight plan); (6) Associated aircraft direct operating costs, according to a general, user-specified function for cost of delay time.
SCHEDULE RELIABILITY	(7) Distribution of arrival times of flights relative to scheduled arrival times; (8) Distribution of arrival times for flights defined to be members of a flight "bank" relative to scheduled arrival times for the bank. (9) Distribution of the percentage of other flights in a bank with which each member of a bank connects.
AIRCRAFT UTILIZATION	(10) Number of aircraft of a given type employed in performing a specified part of a daily schedule of flights.*

* These metrics have not yet been implemented in the ASCENT testbed

Table 5.2.1: ATFM Performance Evaluation Metrics

5.3 THE DRAPER ATFM SIMULATION TESTBED: ASCENT

Draper Lab and MIT have been working together to investigate air traffic flow management concepts since 1989. Since 1991, as part of that collaboration, an ATFM simulation testbed—ASCENT (ATFM System Concept Evaluator for New Technologies)—that has been designed and implemented to evaluate the system-wide impact of new procedures, technologies, and improved infrastructure under existing or anticipated future approaches to ATFM. The current version of ASCENT contains:

i) models for a national network of capacitated[5] and non-capacitated[6] airports;
ii) algorithms for planning ground holds and for allocating mandated delay between the ground and the air;
iii) algorithms for (airline) tactical planning of arrivals at airports
iv) a system level simulation of a day's activities in the National Airspace System (NAS);

5 At a capacitated airport, capacity can be less than demand. Here, ATFM planning deals with all flights into and out of capacitated airports.

6 At a non-capacitated airport, demand is always less than capacity.

v) database and analysis capabilities.

Supporting utility programs include:

vi) models to simulate the evolution of airport weather and capacity;

vii) a tool for generating OAG-like demand schedules at airports.

Figure 5.3-1 illustrates a subset of the many output window formats available from ASCENT's graphical user interface.

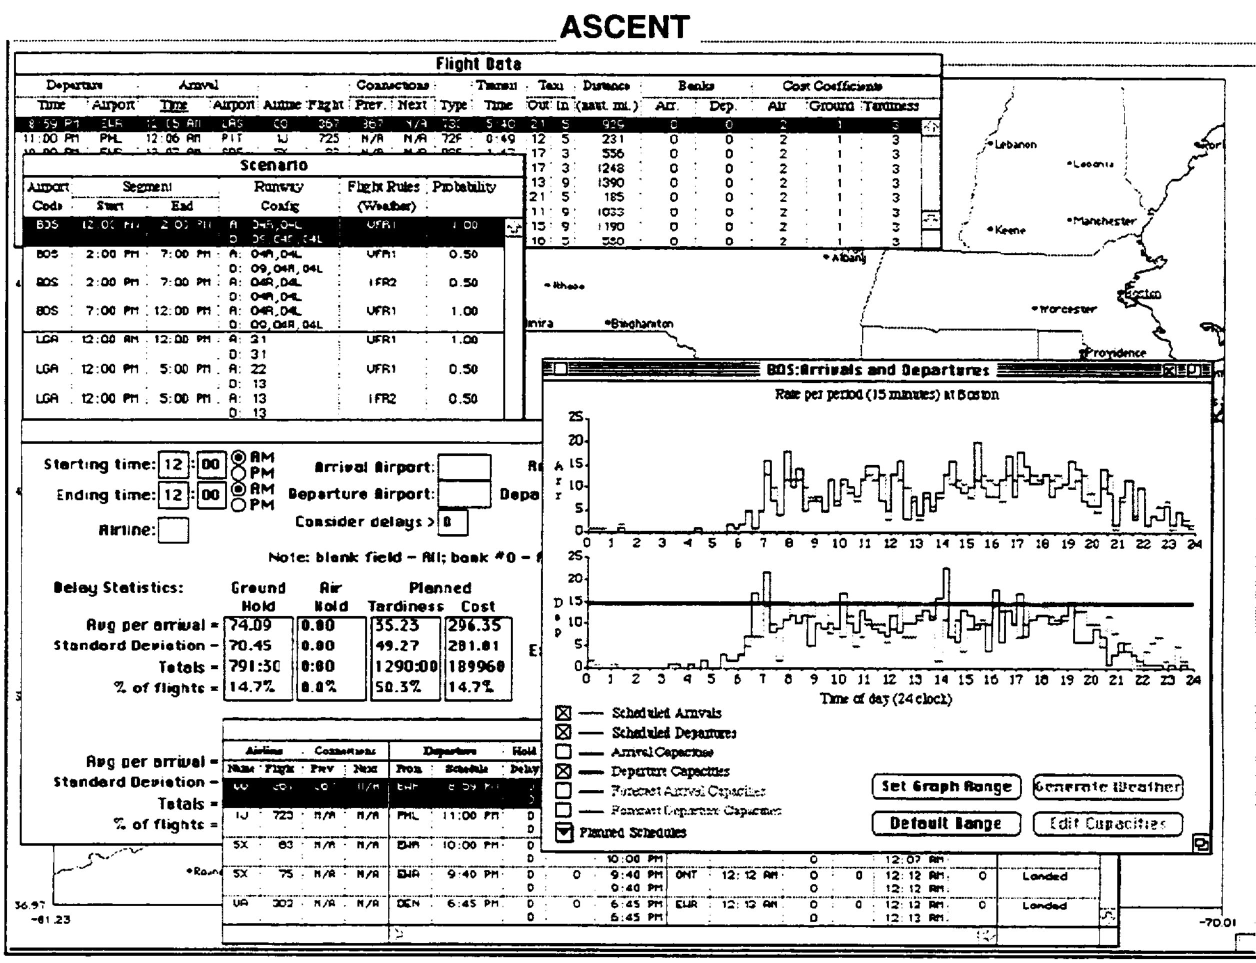

Figure 5.3-1: ASCENT Windows

The map window displays a geographic view of the airports in the flight data, with capacitated airports in red and underlined. Arrival and departure graphs for any capacitated airport can be accessed by pointing and clicking at the airport on the map. In the graph window that opens, the user can select any subset of the following to be displayed as a function of time, for both arrivals and departures: schedules, capacity profiles, forecasts, plans and realizations.

The flight data window offers editing, sorting and visualization of the flight data, a superset of OAG schedules. Capacity scenarios are shown in the scenario window. The bottom window shows the result of a planning and simulation run, with detailed information on a flight-by-flight basis; the user can sort on a number of variables, e.g., air delay, tardiness, arrival time. The last window visible contains summary results. The user can examine summary statistics for a planning and simulation run, and can further filter this by arrival airport, departure airport, airline, time of day, etc.

ASCENT has been designed so that it can be used by a single analyst, requiring a minimum of overhead activity associated with defining and setting up scenarios and performing analyses. It is capable of evaluating candidate air traffic flow management approaches across a spectrum of scenario variations. Flight schedules (demand) and airport capacities (supply) have been determined to be the most significant defining factors for any given scenario. Tools have been created to allow user interaction in the creation of each of these scenario components.

Through the demand-generation tool POAGG (see Section 5.5.1), the user can easily generate OAG-like hypothetical flight schedules for a network of airports. POAGG uses a combination of heuristics and mathematical programming to create statistically realistic flight schedules that satisfy user-specified parameters. These input parameters include for each airport:

(a) The number and hourly distribution of arrivals.
(b) The percentage of flights that connect to each of the other airports in the network.
(c) Directional travel times between airport pairs.
(c) The presence, if any, of shuttle flights and their characteristics.
(d) The presence, if any, of airline flight "banks" and their characteristics.

Once defined, a parameter set can be saved and edited to create new scenarios.

Airport capacity can be defined explicitly by the user, or can be generated automatically by a saw-tooth wave model of weather (Yu, 1996) that has realistic spatial and temporal correlation characteristics. The airport arrival and departure capacities are modeled using the FAA Engineered Performance Standards.

In setting up a simulated test case, the analyst selects a flight schedule and an airport capacity scenario as inputs. One of a set of ground-holding/arrival slot allocation algorithms is selected to create planned aircraft ground holds and slot allocations for the day. Reductions in en route times due to free flight, reductions in airport ground delay times due to the improved ground traffic management or increases in effective airport capacity due to improved arrival sequencing due to, for instance, CTAS (Erzberger, 1993) can also be selected or specified by the analyst. Once the test case set up is completed, the simulation of a day in the NAS is realized, and the resulting delays and other desired evaluation metrics are computed. Note that when weather/capacity are modeled probabilistically, their realizations may not exactly match forecasts that may have been used by algorithms that plan for ATFM activities. If at some point during the simulated day, a (weather or capacity) forecast changes, the analyst can choose to exercise an algorithm to replan ground holds or select an algorithm to tactically resequence arrivals at a given airport, both on the basis of the current state of the system and the new forecast. The analyst can also run an N-day Monte Carlo simulation based on probabilistic capacity scenarios and travel times.

5.4 Models and Algorithms

Modeling airline behavior and quantifying the benefits of alternative ATFM concepts are the most challenging aspects of evaluating decentralized AFTM.

5.4.1. Alternative Concepts for Allocation/Assignment of Arrival Slots

This section discusses columns 1-3 of Table 4.2-1.

5.4.1.1. Modeling Ground-Hold Planning

Ground-holding is the most critical and, if done well, the most effective device that can be employed by ATFM whenever delays are anticipated to be severe. This is particularly true in light of the general desire in ATM to keep airborne delays to a minimum. Achieving the correct mix of ground and airborne delays through ATFM may result in large cost savings to airlines and their passengers. In 1994, ground delay "programs" were applied by Central Flow Control at the ATCSCC (Air Traffic Control System Control Center) on 1089 occasions, resulting in about 400,000 hours of aircraft ground delay with an estimated direct operating cost to the airlines of $625 million. Particular emphasis will therefore be placed here on a detailed examination of ground-holding and of decision-support systems for this purpose.

Ground-holding is typically imposed on aircraft departing for congested airports or scheduled to traverse congested airspace. The motivation is that, as long as a delay is unavoidable, it is safer and less costly for the flight to absorb this delay on the ground before take-off, rather than in the air. Unfortunately, deciding how much ground-holding delay to assign to a flight is far from simple (Odoni, 1987; Andrews, 1993). The reason is that *it is difficult to predict how much delay a flight will actually suffer*, because sector capacities and, especially, airport capacities are often highly variable and can change quickly over time, as weather changes or other events occur.

It is an often insufficiently appreciated fact that airport capacity is a *random variable*. At any given airport, capacity during any period of a day can assume one of several different values, depending on weather conditions (visibility, cloud ceiling, winds, precipitation), the mix of aircraft types, the mix of operations (arrivals vs. departures), the runways in use, equipment outages, human factors, etc. It is not unusual in the United States to encounter a 2:1 or even 3:1 ratio between the highest and lowest capacities of any given airport.

Small changes in visibility at ground level or in the cloud-cover may translate into large differences in airport capacity. Meteorologists cannot yet reliably predict such changes to this level of accuracy, even over a very short time-horizon of an hour or less. Thus, ground-holding decisions must be made under uncertainty and must consider the trade-off between "conservative" strategies that may at times assign excessive ground-holds and more "liberal" ones that may result in more expensive airborne delays. Airlines in the United States contend that TFM tactics have tended to err excessively on the conservative side, i.e., that there may be too many instances in which airport landing capacity (at destination airports) is being wasted, while aircraft stand waiting on the ground at the airports of origin. The MAR (Managed Arrival Reservoir) program has been instituted, in part, to address this concern.

A set of algorithms for modeling ground-holding decisions has been implemented within ASCENT, each representing a different stage of development in TFM automation support. These algorithms model:

(a) *Passive*, representing a system in which ground-holding is not used (aircraft are allowed to take-off when ready for departure);

(b) *First-Scheduled, First-Served (FSFS)*, with and without MAR, approximately representing current TFM practice (FSFS looks at one airport at a time and does not consider the network-wide effects of "local" ground-holding decisions); and

(c) *Collaborative Slot Assignment (CSA)*, FSFS with airlines having the authority to utilize their allocated slots as they wish.

Other ground hold algorithms have been implemented, but were not used for the evaluations presented in this report.

Simulations with these three algorithms will thus provide approximate estimates of some of the metrics in going from "no ATFM system" (*Passive*) to the existing system (*FSFS*) to the viable medium-term scenario (*CSA*).

5.4.1.2. Arrival Slot Scheduling

A variety of algorithms have been developed to minimize total weighted flight tardiness, with the weights chosen to represent the business objectives of the airlines. Bertsimas and Stock (1994) developed an optimal formulation. Brunetta et al (1995) developed a maximum marginal return sequencing heuristic that can be used strategically for both slot allocation and assignment as well as tactically for arrival sequencing. In addition, the arrival sequencing heuristic lends itself to collaborative use by the FAA and airlines. The FAA can first run FSFS on the total schedule and total airport capacities. The slots allocated to each airline are then used by each airline to sequence its own flights using the heuristic. This is what is called Collaborative Slot Assignment (CSA) in this report.

The heuristic is fast, solving problems with thousands of flights in under a second of CPU time on a low-end workstation. It can be used tactically as needed during the day as the forecast of future system demand and capacity changes. It has provided excellent plans within a few percent of the optimal solution in the scenarios tested. It is flexible, based on priority functions that can be dynamically changed by the user (e.g., airline). Each airline can utilize priorities that best suit its operations (e.g., high priorities for certain origin-destination pairs, higher priorities for flights with higher passenger loads, higher priorities for aircraft that have important connections, higher priorities for keeping banks together). Each airline can change its own priorities dynamically as a function of the state of system and the airline's business objectives.

5.4.1.3. Bank Preservation

Airlines that operate in hub-and-spoke environments also have the additional problem of trying to maintain the integrity of their banks of flights. By allowing airlines to allocate their own slots, airlines will be able to accommodate their banks better. For example, a bank may have originally been scheduled to arrive over the course of 30-40 minutes, but may be assigned a set of arrival slots which span several hours. The airline's response may be to choose a particular subset of flights of the bank and keep them together, delaying or canceling the rest. Another airline facing the same problem may choose differently. This represents an additional capability for advanced airline collaboration.

A model developed by Milner (1995) takes into account the dependencies of flights in a bank. Specifically, the model includes information regarding the total delay that flights experience when they arrive in a bank of flights. This delay is greater than the delay experienced by individual flights because of the time flights will spend at the hub airport waiting for the completion of the bank. An airline attempting to allocate its arrival slots would either assign a flight to arrive as part of the bank in which it was scheduled to arrive, assign the flight to a slot after the bank was completed, or cancel the flight outright. If a flight were assigned to a slot after its bank's completion, passengers on that flight would likely miss their connections. Furthermore, passengers at the hub airport connecting onto the next expected flight for that particular aircraft would also be delayed. The assumed objective of the airline in the model is to minimize a weighted combination of the delay incurred by flights that remain with their banks, the delay incurred by flights separated from their banks and the cost of canceling flights. More details and results can be found in Section 5.5.3. This models another component of advanced airlines collaboration and focuses on collaborative decision-making that impacts bank operations.

5.4.2. Alternative Concepts for En Route Control

This section discusses column 4 of Table 4.2-1.

One of the principal benefits expected from free flight concepts is a reduction of en route flight times as a result of the ability to operate with more direct and wind-optimized flight trajectories. The existing version of ASCENT does not simulate explicitly the en route segment of flight. It is possible, however, to evaluate the benefits of reduced flight times to the airlines by varying systematically the assumed en route flight times for each origin-destination pair in the simulator and observing the effects of flight-time reductions on delays and schedule reliability. In other words, *ceteris paribus* (i.e., for a given TFM system) we wish to observe the impacts of an $X\%$ (or an X minute) reduction in flight times on "downstream" delay propagation through a set of scheduled flights, with X being varied systematically (possibly taking into consideration the length of the flight and type of aircraft involved). Increased schedule reliability has important consequences for aircraft productivity. Large improvements in reliability would make it possible for airlines to fly a larger number of flights with the same number of aircraft. Ultimately, to realize the full benefits of reduced flight times, those reduced times must be reflected in reduced scheduled times (i.e., OAG times) which, in turn, lead to the potential for improved fleet utilization.

5.4.3. Alternative Concepts for Terminal Area and Ground Movement

This section discusses columns 5 of Table 4.2-1.

Modest airport capacity increases during the next ten years will be made possible through such programs as CTAS (more efficient sequencing and spacing in the near-terminal area), the Surface Movement Advisor (SMA) program, Airborne Information for Lateral Spacing (AILS), and several other terminal area innovations (most of which are being investigated under NASA's ongoing Terminal Area Productivity program). ASCENT models some of the system-wide effects of increased airport capacity on TFM-related delays and other costs by varying capacities parametrically at selected airports. For a given TFM system, the differences in the values of the various evaluation metrics before and after the capacity increases realized by the introduction of these effective capacity increasing procedures and technologies will be used to assess the magnitude of the resulting benefits.

5.5 SUPPORTING RESEARCH

5.5.1. Airport Demand Schedule Modeling

This section describes an algorithm previously developed, implemented and tested under Draper Lab internal funding and enhanced under this project for generating hypothetical, OAG-like airport demand schedules. The two major enhancements have been porting the application to platform-independent code and improving the fidelity of the departure model. The objective of the algorithm is to produce daily airport schedules that adhere to flight connectivity and arrival rate parameters for a set of user-specified capacitated airports, so that the flight schedules produced by the algorithm are statistically similar to actual or anticipated OAG schedules and thus are useful for ATFM analysis. It is important to note that the generated schedules are not designed to be implemented by airlines. The algorithm is implemented in a computer application known as the Pseudo-OAG Generator or POAGG. The POAGG algorithm extends previous flight schedule generation research (Hocker, 1994).

5.5.1.1. Motivation and Objectives

The development of a schedule generator was motivated by the need for airport demand schedules to support air traffic management research. The decision to model commercial demand was based on observations of the dominance of commercial traffic at major airports over other forms of air traffic. In addition, the current practices within the FAA of forecasting demand primarily involve commercial traffic.

The current source of demand schedules for ATM analysis is OAG flight schedules. There are limitations in how much one can use OAG schedules, since they represent existing or historic air traffic demand. Given the predicted, ten-year 4% annual growth rate in air traffic demand levels[7], OAG schedules are insufficient for a simulation of the future of the air traffic environment. Furthermore, future changes in airline scheduling practices are not reflected in current or past OAG schedules. By offering the capability to generate hypothetical schedule scenarios that capture the growth in air traffic demand and the dynamic nature of airline

[7] FAA annual forecast for domestic traffic released April 1993.

scheduling, POAGG schedules are a necessary supplement to OAG schedules as demand input for ATFM analysis.

As a secondary motivation, hypothetical schedules may be useful in the analysis of proposed air traffic capacity infrastructure enhancements. Given the long term and costly effort required to substantially increase an airport's capacity, e.g., building a new runway, it is appropriate to measure the cost effectiveness of an enhancement against future demand levels.

The basic characteristics of commercial air traffic schedules modeled in the POAGG algorithm are flight connections, hubbing operations, and shuttle services. These characteristics are the principal factors that both influence and are significantly affected by ATFM actions and system delays. The modeling effort has been intended to produce schedules that preserve these characteristics, in such a manner that hypothetical schedules generated by the algorithm are statistically indistinguishable from real OAG schedules, for the same airports and similar demand levels. Information such as airline, flight number, aircraft type, etc. for the hypothetical flight itineraries is modeled to enhance schedule realism, but is not intended to model actual airline scheduling practices at a high fidelity.

- *Connecting Flights*

Airlines routinely schedule an aircraft to perform multiple flight legs throughout the day, creating an interdependency between an incoming flight and the departing flight of the same aircraft. The departing flight is referred to as the *connecting flight*. *Layover time* refers to the scheduled time interval between the arrival time of an incoming flight and departure time of its connecting flight. Due to flight interdependencies caused by connections, congestion delays often propagate to connecting flights throughout the day. In the same manner, flow management actions by the FAA, may also influence connecting flights. An assigned ground delay on one flight may result in derivative delays on subsequent flights of the same aircraft.

- *Hubbing Operations*

The ripple effect of delays due to flight interdependencies is exacerbated by the airline hubbing system. A hub airport is an airline's central operating location within a geographical region. Major airlines have 3 to 4 hubs. At a hub airport, an airline will typically schedule 20 to 40 incoming flights within a specific interval of time. This grouping of flights constitutes an *arrival bank*. Following a sufficient period of time to allow for passenger connections and aircraft preparation, the airline schedules a *departure bank*. This scheduling practice is an efficient use of available aircraft and creates multiple connection opportunities for passengers traveling in the coupled arrival and departure banks.

However, banks create many flight interdependencies. Departure bank flights mostly consist of connecting flights from an associated arrival bank. As a result, airline hubbing systems are sensitive to air traffic delays. A delay of a single flight in an arrival bank, may cause an airline to delay the departures of several flights in the departure bank to accommodate incoming passengers connecting to multiple outgoing flights.

- ***Shuttle Services***

A shuttle service is established by an airline to provide air travel between an airport pair on a regular basis throughout the day. Typically, shuttle flights depart on an hourly basis. Aircraft within a shuttle service can be dedicated to travel back and forth between the airport pair.

Three observed characteristics of airline schedules that are not in the POAGG model, per se, are aircraft-specific flight and layover times, and departure time biases.

Actual flight times between two given airports differ for jet engine and propeller-driven aircraft. Accordingly, airlines schedule a flight's arrival time to reflect aircraft type. In POAGG, the model assumes a uniform fleet on each route, i.e., scheduled flight times are assumed to be same for all flights traveling between two given airports.

This assumption does not affect the validity of the model, because in practice, most routes are scheduled with similar aircraft types. The uniform fleet assumption is largely correct except for some short-range routes, in which propeller-driven aircraft and jet engine aircraft may be scheduled between the same airport pair. POAGG does model the difference in a flight's direction for a given airport pair, e.g., the flight time from Airport 1 to Airport 2 is modeled separately from the flight time from Airport 2 to Airport 1. This reflects the difference in East-West and West-East flight times between the same airports due to jet stream winds.

In practice, airlines schedule layover times according to aircraft type. The type and respective size of an aircraft influences the time required to prepare for connecting flights and transfer passengers, i.e., the *turnaround* time. Furthermore, certain airlines have faster turnaround times than others for the same aircraft type. Although POAGG models layover times as random variables, the variation does not reflect aircraft types or airlines. Correct layover times would enhance the realism of generated schedules; however for ATFM analysis, precise layover times are not critical.

Departure time biases refers to the practice of favoring departures on the hour and half-hour. To a lesser degree, airline schedules contain departures on the five minute interval, e.g., 8:05, 8:10, 8:15 etc., over other minutes within the hour. Bias toward departures on the hour and half-hour is indirectly simulated in the shuttle service model.

5.5.1.2. Implementation of the Model

ASCENT will accept either modified OAG flight schedules or POAGG generated schedules as air traffic demand input. POAGG has been ported to platform-independent C, using text input and output files. The components of the flow chart in Figure 5.5.1.2-1 represent different phases of the schedule-generation process.

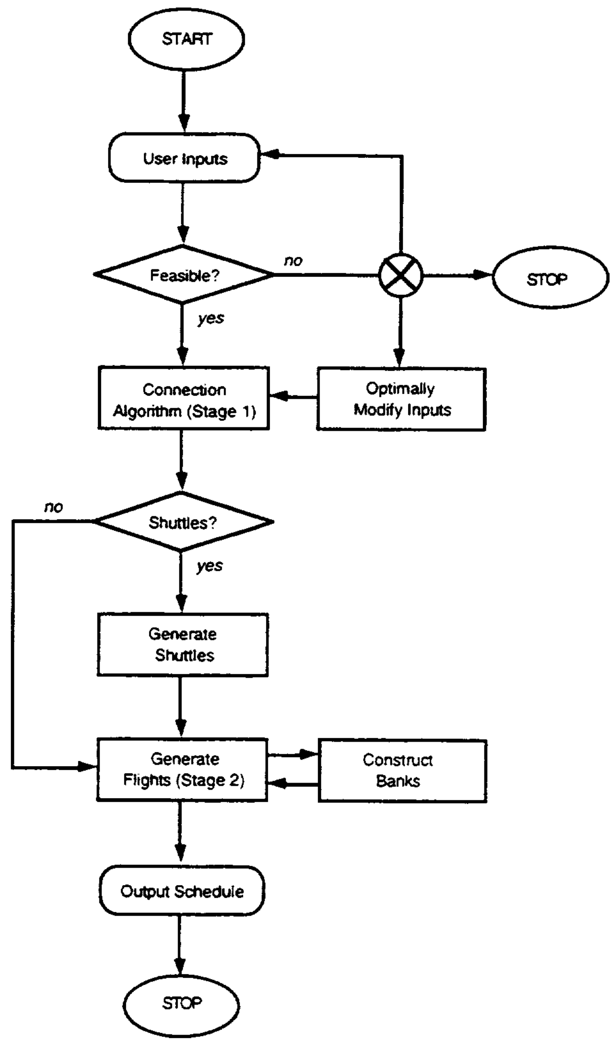

Figure 5.5.1.2-1: POAGG Flow Chart

5.5.1.3. POAGG Functionality

POAGG is a user-driven application. The user specifies a network of airports by defining parameters that describe each airport. The following are lists of input and output data for POAGG.

The general inputs are:

1. number of capacitated airports
2. number of non-capacitated airports
3. average flight time to each non-capacitated airport
4. period length
5. cutoff time for connecting flights to non-capacitated airports
6. probability mass functions (PMFs) for flight times to non-capacitated airports
7. minimum required layover time
8. maximum allowable layover time
9. upper parameter for layover probability density function

For every capacitated airport, the inputs are:

1. three-letter airport name, e.g., JFK
2. total scheduled arrivals
3. relative frequency distribution of arrivals
4. relative frequency of connecting flights to other capacitated airports
5. PMF and carrier names for airlines operating at the airport
6. PMFs for airline fleets

Input data describing the relationships between airport pairs are:

1. directed scheduled flight times
2. connection percentages

Shuttle input data are:

1. airport pair
2. carrier name
3. aircraft type
4. start time
5. stop time
6. interval between shuttle flight departures

Arrival/Departure banks input data are:

1. hub airport name
2. arrival bank start time
3. arrival bank stop time
4. minimum objective time
5. number of flights in arrival bank

Output consists of flight records that include

1. departure and arrival time
2. airline, aircraft type
3. previous and next flight of aircraft
4. arrival and departure bank membership

During the original development of POAGG, a number of validation tests were performed. The primary change to POAGG's functionality in this project has been improved modeling of departures. This was done by enforcing a set of constraints on when aircraft could start their itinerary.

Tests were performed comparing the enhanced POAGG against OAG data. POAGG was run on input data taken from an analysis of an OAG schedule. Table 5.5.1.3-1 compares the number of POAGG departures from 3:00 PM to midnight from five airports (BOS, DCA, EWR, LGA and PIT) to the corresponding number of departures found in the OAG schedule.

Time	BOS		DCA		EWR		LGA		PIT		Total	
period	OAG	POAGG	OAG	POAGG	OAG	POAGG	OAG	POAGG	OAG	POAGG	OAG	POAGG
15:00-	21	27	9	15	14	8	17	16	22	23	83	89
15:30-	27	17	12	13	14	19	15	10	19	22	87	81
16:00-	29	30	11	10	30	27	16	13	7	25	93	105
16:30-	16	21	12	12	11	16	16	17	38	21	93	87
17:00-	31	23	18	13	19	21	14	19	11	14	93	90
17:30-	23	28	10	16	32	28	18	14	26	24	109	110
18:00-	27	29	13	10	27	20	20	21	32	24	119	104
18:30-	21	17	14	14	18	15	14	15	8	12	75	73
19:00-	25	24	17	15	7	24	13	17	2	1	64	81
19:30-	20	25	8	13	37	23	19	14	6	6	90	81
20:00-	20	26	18	12	20	22	13	17	28	34	99	111
20:30-	11	16	4	12	7	11	12	13	24	32	58	84
21:00-	9	17	8	14	2	15	8	17	1	20	28	83
21:30-	2	3	2	2	3	2	1	1	46	0	54	8
22:00-	12	2	1	2	2	1	2	1	13	1	30	7
22:30-	5	0	0	1	3	0	0	4	3	2	11	7
23:00-	1	0	0	0	1	2	0	2	1	0	3	4
23:30-	3	0	0	1	2	0	0	0	2	0	7	1
Total	654	641	373	387	577	569	481	462	565	553	2650	2612

Table 5.5.1.3-1 Departures

Table 5.5.1.3-2 compares the number of connections made within the five airport network; note that airport XXX represents all airports that have flights into or out of the five airport network. Arrival distributions were compared but are not presented since there was an exact match between POAGG and the OAG schedule.

OAG

	BOS	DCA	EWR	LGA	PIT	XXX	Depart.
BOS	0	32	35	34	9	544	654
DCA	31	0	18	31	8	285	373
EWR	35	17	0	0	14	511	577
LGA	34	31	0	0	9	407	481
PIT	9	8	14	8	0	526	565
XXX	542	286	515	407	525	0	2275
Arriv.	651	374	582	480	565	2273	4925

POAGG

	BOS	DCA	EWR	LGA	PIT	XXX	Depart.
BOS	0	30	29	29	9	544	641
DCA	30	0	17	30	8	302	387
EWR	40	10	0	0	12	507	569
LGA	34	30	0	0	8	390	462
PIT	10	6	15	7	0	515	553
XXX	540	302	522	414	531	0	2309
Arriv.	654	378	583	480	568	2258	4921

Table 5.5.1.3-2: Connections

5.5.1.4. Future Extensions

The airport demand model presented in this section offers the capability to generate OAG-like schedules for multiple airports with variable connection percentages. The current implementation of the model in the POAGG application is fast and robust. POAGG has been delivered to multiple ATFM researchers as a support tool. It has also been used extensively within this project.

The numerical method implemented in POAGG effectively models commercial demand for a network of specified airports. The method can potentially be extended to model general aviation (GA) and military air traffic. One alternative is to insert GA and military flights into the commercial schedule using a Monte Carlo method. GA and military flights could randomly be assigned arrival times into specified airports based upon historical or hypothetical distributions of GA and military airport demand.

Additionally, the POAGG numerical method can be extended to include a higher fidelity model of individual airline scheduling. Enhancements could include modeling flight times as airline and aircraft specific, and biasing schedule departure times based on airline preferences. The process of assigning connecting flights could be modified to model connections at the airline level. This would require input parameters and constraints that describe the scheduling practices of the individual airlines.

5.5.2 ASCENT Enhancements

The core functionality of the Draper ASCENT testbed was developed under Draper Lab IR&D funding from 1991- 1995. During the current project a number of enhancements to the testbed were developed to provide the capabilities needed to perform analyses required to support investigations of ATFM concepts described herein. A brief description of the major enhancements follows.

5.5.2.1 National network

The ASCENT testbed's design supports the full NAS network of airports, with provision for up to forty[8] capacitated airports. Data associated with each capacitated airport is input from an airport data file. All flights arriving or departing at any of the capacitated airports are planned and simulated in the testbed. Further, scheduled, planned and actual plane arrival/departure rates and scheduled, forecast and actual airport capacities are maintained and may be displayed for each airport. Planned and actual plane arrival/departure rates are maintained over multiple scenarios while the testbed is an active application.

In order to gauge the capability of ASCENT to handle the entire set of one day's flights in the NAS, it was tested on files of approximately 5K, 15K and 25K flights . The run-times on a mid-range PowerMac (120 MHz 604) for one day's planning, simulation and evaluation are given in Table 5.5.2.1-1. Pre-processing time includes initial setup of connections of flights and sorting for use by planning algorithms and the simulation. The 25K file is the largest flight data file available to us. Based on the results below, it is likely that (with enough computer memory) ASCENT can handle an entire day's flights for the NAS (approximately 40-50K).

Number of flights	Number of capacitated airports	Pre-processing time	FSFS Planning run-time	CSA Planning run-time	Simulation run-time
4,925	38	3 seconds	1 second	2 seconds	10 seconds
14,765	38	37 seconds	5 seconds	5 seconds	28 seconds
25,102	38	94 seconds	16 seconds	16 seconds	52 seconds

Table 5.5.2.1-1 Single Day Run-Times for ASCENT

[8] Forty was chosen because we had access to capacity data for only 38 airports. There is no inherent limit on the number of capacitated airports.

Figure 5.5.2.1-1 illustrates the total run-time needed for one day's planning with the FSFS algorithm, simulation and evaluation for the three cases.

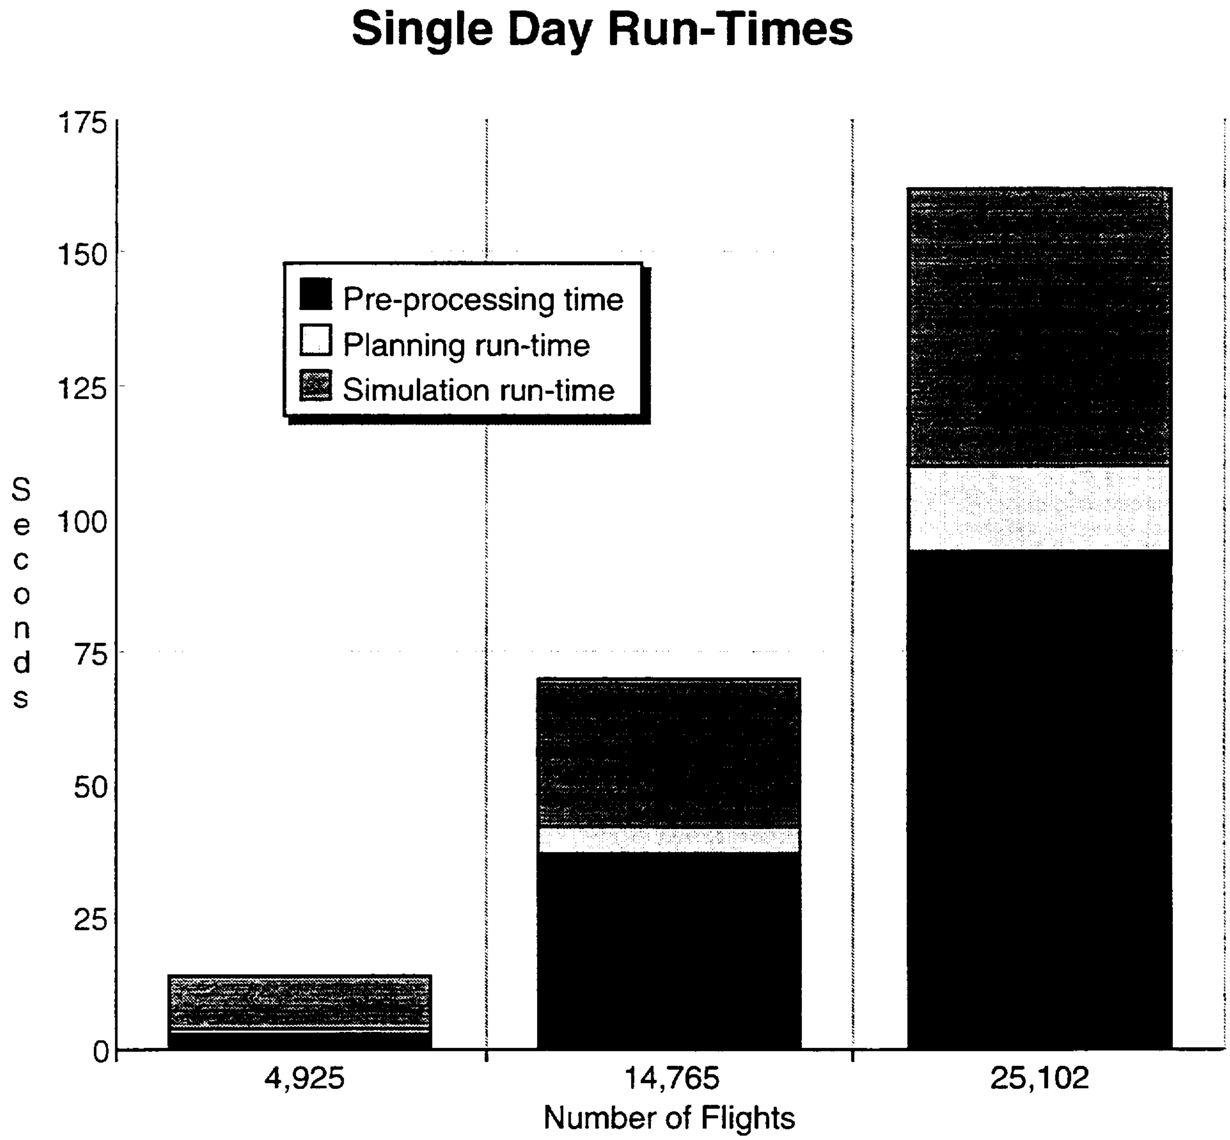

Figure 5.5.2.1-1: Single Day Run-Times

5.5.2.2 Hub airports and banks

ASCENT supports identification of arrival and departure banks. As with other flight information, the bank information is displayed in the Flights window and may be used for sorting and filtering for output. The information is also available to all of the planning algorithms as required.

5.5.2.3 Updated FAA model of Ground Delay Program-MAR, Local

The model of the FAA's Ground Delay Program uses single airport First-Scheduled First-Served planning logic. In addition, it now includes the ability to model local Ground Delay Programs by specification of a distance over which flights are not given planned ground delays.

The Managed Arrival Reservoir (MAR) is also modeled. MAR results in ground holds being decreased by a maximum of a user-supplied time, with a default of 15 minutes. This allows the planes affected to take advantage of any available slack time in the system (due, for example, to statistical deviations in the flight times of planes or variations in arrival capacity).

5.5.2.4 Models of CTAS, UPR

The Center TRACON Automation System (CTAS) is currently modeled in the ASCENT testbed as a user-set percentage increase in airport capacity for the airports that are CTAS-equipped. User Preferred Routing (UPR) is modeled by allowing user-set variation (typically reduction) of individual flight transit times, independent of the actual scheduled departure and arrival times.

5.5.2.5 Collaborative Slot Assignment (CSA) Algorithm

The algorithm for modelling CSA that is described in Section 5.5.5 has been implemented and integrated into ASCENT.

5.5.2.6 Arrival and Departure Capacities

Support for modeling departure capacities at the simulated airports has been added to the testbed. Currently, when airport departure capacities are exceeded, planes waiting to depart will incur taxi-out delays (see below). Alternatively, it would be possible to model departure capacity overflow as gate delays, but this would affect incoming flights due to the unavailability of gates. More work needs to be done in this area to accurately model the current operational practices of the airports. Displays of planned and actual arrivals and departures as well as actual available arrival and departure capacities are provided for each simulated airport. As with arrival capacities, the effects of weather on departure capacities are simulated in accordance with the flight rules and runway configurations for each of the simulated airports.

5.5.2.7 Taxi-in and taxi-out

The ASCENT testbed now models the normal taxi-in and taxi-out times for each flight, as well as taxi delays incurred due to limits on the airport arrival and departure capacities. Currently, this is a low-fidelity model that increases taxi delays as more planes are queued for arrival or departure, i.e., delays are increased as a function of measure taxi traffic congestion.

5.5.2.8 Transit times and time zones

In order to more accurately model flight times, and also to support User Preferred Routing (see above), flights are simulated using an expected transit time (time from wheels-off to touchdown) that is input along with the other flight data. This is an enhancement over the previous use of schedule times to infer the transit time for a flight. Planning may be done using a choice between the expected transit time and the inferred transit time based on schedule.

In order to reflect the effects of travel across time zones, planning and simulation are now performed using Universal time. For display purposes, times are still adjusted to local time.

Using Universal time for simulation increases the fidelity by avoiding possible errors due to network effects caused by connecting flights across time zones.

5.5.2.9 Cost functions

The testbed supports a three-coefficient cost function for each flight that may be used for planning as appropriate. In particular, each flight has an input set of three integer coefficients that define that flight's cost function. Currently, a linear cost function is applied using the coefficients as weights for ground delay, air delay, and total tardiness, summed to form the total cost. For future work, the cost function could correspond to individual airline supplied or hypothetical airline utility functions.

5.5.2.10 Improved forecasting model

Forecasts of airport capacities for the purposes of planning can be statistically linked to the actual capacities that are realized during simulation. There are 5 types of forecasts as follows: 1) the forecast is equal to the realization, 2) both forecast and realization are simulated using the same distribution, 3) the forecast is simulated using a uniform distribution, 4) the forecast is "optimistic," the maximum achievable for each period, and 5) the forecast is "pessimistic," the minimum achievable for each period. This models a range of forecasting capabilities by the planner. Since airport capacities can now be simulated stochastically based on input capacity scenarios (see below), this linked forecasting capability is essential to accurately modeling and evaluating future planning efforts.

5.5.2.11 Improved reporting capabilities

All window displays in the testbed have been upgraded to display information associated with the above capabilities. This includes upgraded displays for each airport to reflect departure as well as arrival information and also upgraded summary and histogram displays. In addition, all flight simulation results are output as text files formatted to facilitate post-processing of the data. For the results generated for this effort and presented in Section 6, a FileMaker Pro database application (integrated with Matlab for charting) was designed to post-process the data.

5.5.2.12 Improved representation of capacity scenarios

Airport capacities are now determined based on an input scenario file. The scenario file describes the possible weather (flight rules) and associated runway configurations and the probability distributions of each by individual airport for all times during the day. Each weather/runway configuration corresponds to an arrival/departure capacity at a given airport. The table of capacities and their probabilities is input along with the other airport data from the airport file. At simulation time, the probability distributions are used to select an actual weather scenario (i.e., capacity) for each airport to be used for the simulation. Depending on forecasting quality that was selected for planning purposes, the planner will use the scenarios to do its forecasting of capacities as well.

5.5.2.13 Monte Carlo simulation capabilities

The ASCENT testbed can perform and collect data across multiple simulation runs. In addition to stochastically varying the airport capacities and planning forecasts of capacity (see above), the actual transit times of individual flights can be stochastically varied. By performing multiple runs and allowing desired stochastic variations to occur, a Monte Carlo simulation can approximate long-term average effects for the candidate ATFM concept being evaluated.

5.5.3 Bank Preservation

Modeling airline behavior and quantifying the benefits of alternative ATFM concepts are the most challenging aspects of evaluating decentralized AFTM. In this section, we describe work that focused on quantifying the benefits that may accrue from the airlines' improved ability under decentralized systems to preserve the "integrity" of their flight schedules.

The hub system, adopted by the major carriers in the 1980s, further complicates the effect of flight delays. Under the hub system, an airline schedules many flights to arrive into a hub airport within a short interval of time. Such a group of flights is referred to as an *arrival bank*. Following a short time period to allow for passenger connections and aircraft preparation, the airline schedules a *departure bank* of flights. A major weakness in the hub system is its sensitivity to flight delays. Specifically, ground delays or airborne holds imposed on one or two flights in an arrival bank, may result in airline-imposed delays on multiple flights in the related departure bank. Airline decisions to delay departing flights are based on the percentages of connecting passengers on the late arriving flights and potential downstream effects.

Airlines that operate in hub-and-spoke environments constantly face the problem of trying to maintain the integrity of their banks of flights, whenever major delays occur at these airports. By allowing airlines to utilize their own slots as they deem best, decentralized ATFM concepts may make it possible for the airlines to accommodate their banks better. For example, a bank may have originally been scheduled to arrive over the course of 30-40 minutes, but due to reduced capacity in the system may be assigned a set of arrival slots which span several hours. The airline's response may be to choose a particular subset of flights of the bank and keep them together, delaying or canceling the rest. Another airline facing the same problem may choose differently. What is clear is that an ATFM system operator cannot be as effective as the airlines themselves in making these decisions, because the ATFM operator cannot know perfectly each airline's preferences and utility functions.

Milner (1995) has developed two models which airlines may use in forming their response to arrival slot allocations. The Independent Flights (IF) model describes the problem airlines face in allocating their arrival slots under the assumption that the airline views its flights as being independent of each other. Such a model might be applicable at a spoke airport. He notes that the model is similar to others presented earlier, e.g. Vasquez-Marquez (1991).

The second model takes into account the dependencies among flights in a bank. Specifically, the model includes information regarding the total delay flights experience when they arrive in a bank of flights. This delay is greater than the delay experienced by individual flights because of the time flights will spend at the hub airport waiting for the completion of the bank. In the second model, referred to as the Cancellation/Delay (C/D) model, an airline attempting to allocate its arrival slots would either assign a flight to arrive as part of the bank in which it was scheduled to arrive, assign the flight to a slot after the bank was completed, or cancel the flight outright[9]. If a flight was assigned to a slot after its bank's completion, passengers on that flight would likely miss their connections. Further, passengers at the hub airport connecting onto the next flight that particular aircraft was expected to fly would also be delayed. The assumed objective of the airline in the model was to minimize a weighted combination of the delay incurred by flights that remain with their banks, the delay incurred by flights separated from their banks and the cost of canceling flights.

Milner has shown that airlines can allocate their arrival slots in ways which are consistent with their preferences. In one set of experiments he simulated the arrival of several banks of flights at a hub airport for an airline. Each day was divided into 15 minute periods. Each bank of flights for the airline was scheduled to arrive within a 30 minute interval. A second airline was also scheduled into the airport, the flights for that airline being a constant number per period. This second airline's schedule represented the total schedule for all other airlines operating at the airport. The simulated arrival pattern is displayed in Figure 5.4.1.3-1. In the figure, Airline A is seen to have scheduled banks of four flights arriving each hour, in a 2 period or 30 minute interval. Airline B has four flights arriving every period. In the actual experiments the number of flights in Airline A's banks varied between 20 and 60 flights, while Airline B had a constant 10 flights per 15 minute period.

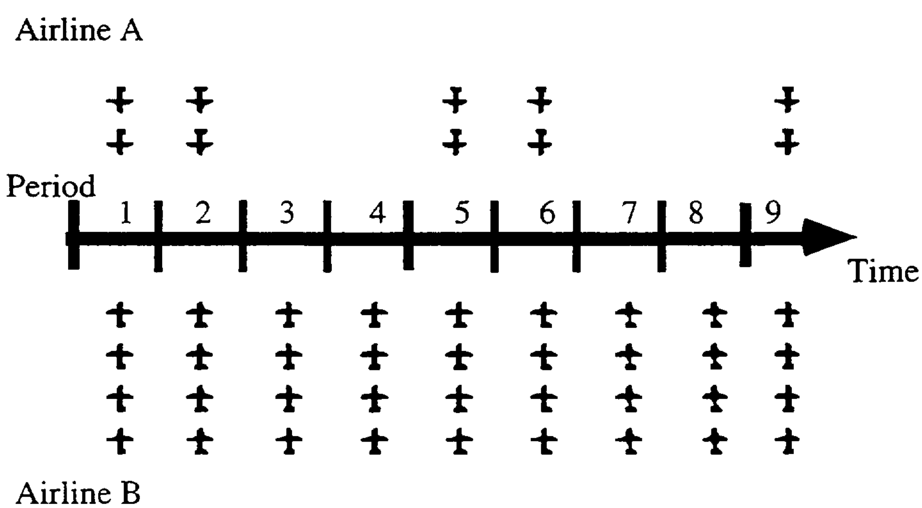

Figure 5.4.1.3-1: Arrival Pattern Used in Experiments

In the experiments, a nominal capacity sufficient to satisfy the entire demand without delays was reduced systematically to 90%, 80%, 70% and 60% of its original amount. Airlines were allocated arrival slots based upon the simulated schedule. Under these conditions, it was shown that an airline could benefit by allocating its arrival slots depending on its preferences. In particular depending on the cost an airline assigns to the cancellation of a flight, the amount of total delay encountered and the number of flights canceled varies.

9 Canceling a flight alleviates the burden of the airlines to feed and house passengers stranded overnight at a hub due to missed connections.

Figure 5.4.1.3-2 displays how an airline might trade off the number of canceled flights and the delay penalty (the total number of periods of delay experienced by flights in the airline's schedule). The trade-off curves are given for banks of 20 through 60 flights. An airline which assigns a low cost to canceling flights would cancel many flights; however, for those flights which were allocated arrival slots, few delays would be incurred. An airline which assigned a high cost to canceling flights would schedule many flights, resulting in high delay costs, but low total cost for the canceled flights. The conclusion from the figure is that different airlines could position themselves at varying places on the curve, a result achievable only with some form of freedom to make this decision as would occur ATFM decentralization.

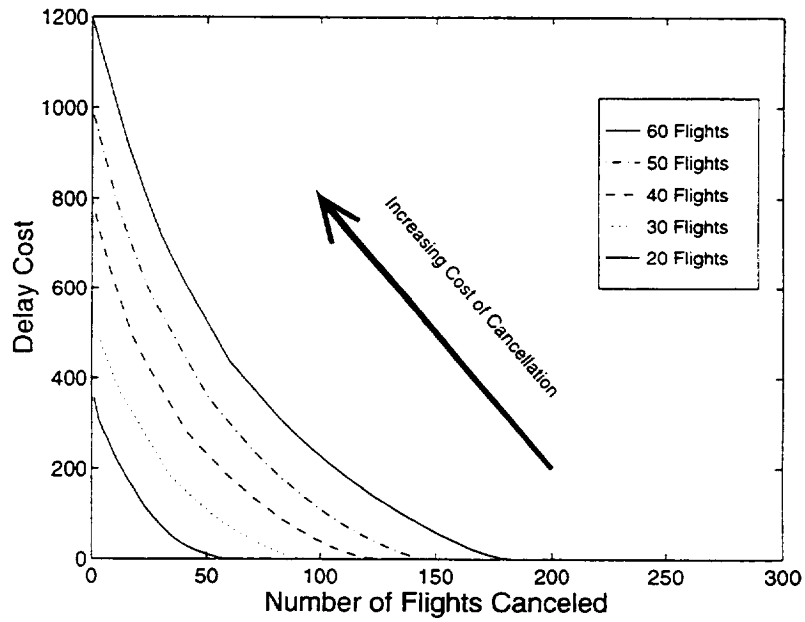

Figure 5.4.1.3-2: Number of Canceled Flights vs. Delay Cost with Changing Cost of Cancellation

These experiments also showed that airlines which consider their flights as being interdependent will allocate their arrival slots in ways which differ from airlines which view their flights as being independent. Figure 5.4.1.3-3 shows that the total amount of delay experienced by airlines which consider their flights as independent is greater than the delay encountered by airlines which allocate arrival slots taking the true dependency of flights into account. The 'x' marks represent an airline using the Independent Flights model, where as the 'o' marks represent the results of the C/D model. The curve in the figure is a best fit for the C/D model's results. The experiments show that when few or many flights were canceled, the total delay is about the same for either model, however, when the total number of flights canceled was between 5% and 20% of all flights, airlines using the C/D model could reduce the delay for each flight flying by 1 to 3 periods (15 minutes to 45 minutes). The IF model selects less expensive flights to cancel, while the C/D model selects flights which will reduce both the delay and cancellation costs. Again, the decision-making freedom available under a partially decentralized system is required for airlines to achieve these results.

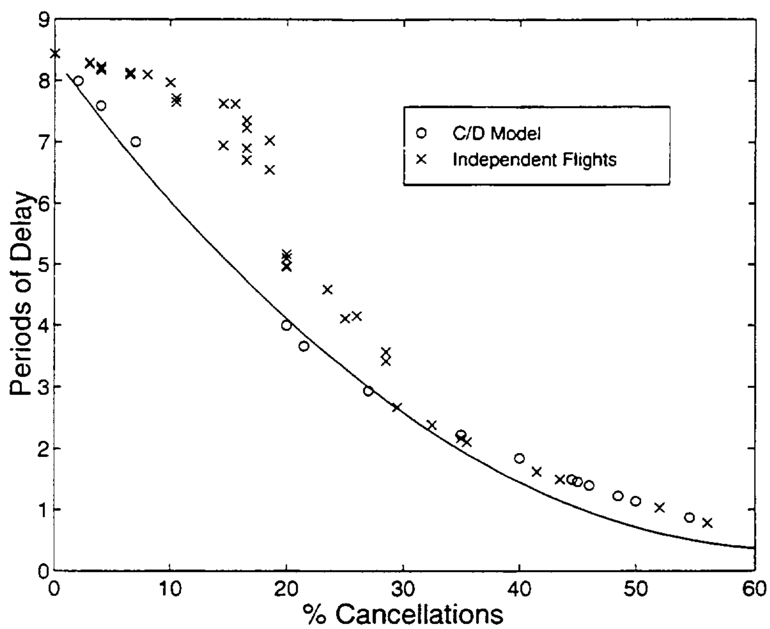

Figure 5.4.1.3-3: Comparison of IF and C/D models

The schedules that are the result of allowing airlines to allocate their arrival slots can vary greatly, depending upon how an airline views the cost of canceling flights and the cost of separating them from their banks. Figure 5.4.1.3-4 shows a result of applying the C/D model for an example where an airline had five banks of flights scheduled to arrive in five consecutive hours. In each fifteen minute period the airline was assigned a varying number of arrival slots based upon the original schedule of all the flights scheduled into the airport. Each distinct slot is depicted by a circle, square or triangle. A circle with a number in it indicates that the slot was assigned to a flight from the indicated bank number and that the flight arrived with the rest of the bank. A square represents a slot not used by the airline because it canceled some flight and did not substitute another into the available slot. A triangle with a number in it indicates some flight from the indicated bank number was assigned to that slot, but the flight was separated from its bank. The figure shows that four flights from bank 2 were separated from the rest of the bank. Even though those four flights arrived immediately after all of the rest of bank 2, the other flights would not wait for the connecting passengers from those four flights. Similarly two flights from bank 3 and one flight from bank 4 were scheduled to arrive after bank 5 was completed, with no possibility for passengers on those flights to make their original connections.

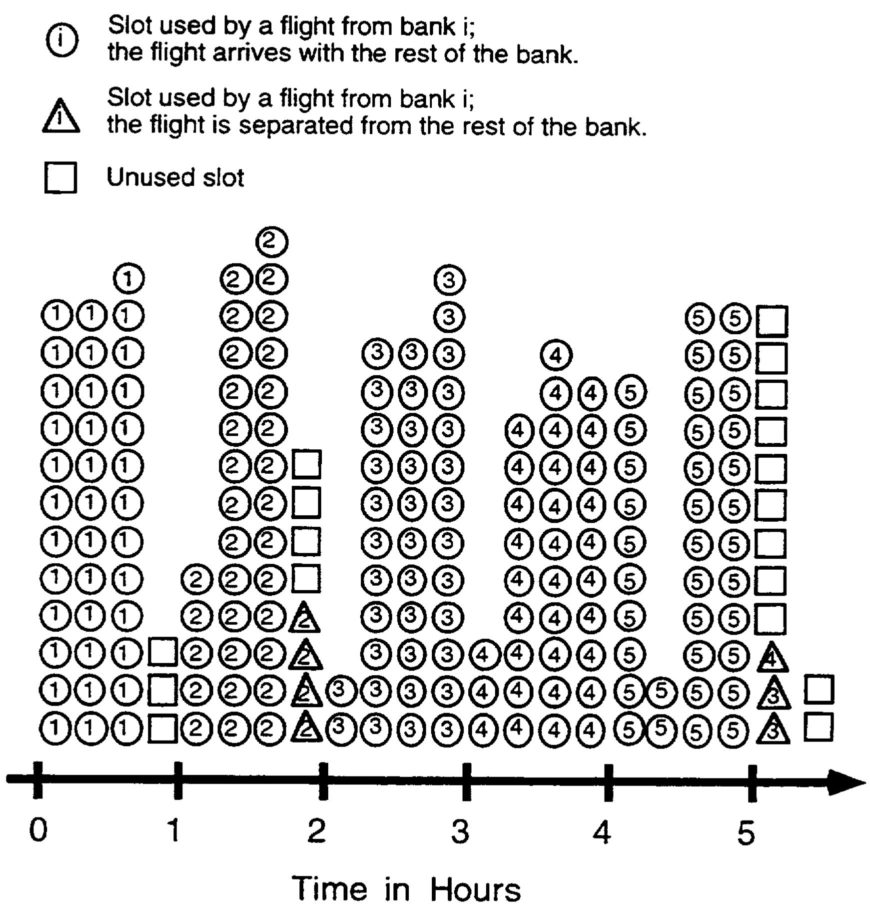

Figure 5.4.1.3-4: Assignment of Flights to Slots by C/D Algorithm

As noted in the example of Figure 5.4.1.3-4, several slots were not used, in particular many of the slots occurring in the sixth hour. Under current practice, airlines have little incentive to reveal the presence of these open slots until it is too late for their competition to occupy them. While allowing a limited form of decentralization can aid airlines in allocating their arrival slots in accordance with their preferences, additional decentralization involving some exchange of slots between airlines as described in row 2d of Table 4.2-1 would be needed so that these unoccupied slots might be filled. Alternatively, either of the more decentralized plans indicated in the row 2e of the table would likely ensure that open slots do not exist during these periods of reduced capacity.

5.5.4 En Route /Terminal Area Transition Region Simulation and Modeling

A part of the effort supported by this project addressed the very important problem of transition between a region of controlled airspace surrounding an airport or other congested area, and a region of uncontrolled airspace employing, for example, the free flight concept of user-preferred routing.

An important question that must be answered is how air traffic can be made to transition safely and efficiently between uncontrolled and controlled airspace, especially under conditions wherein the demand for use of the controlled airspace exceeds its safe capacity. Analysis of this problem yields insight into the additional questions of how large the controlled airspace region must be, and whether its size should be fixed or should be adjusted dynamically to accommodate changing demand conditions. Efficiency concerns would push the controlled region to be as small as possible, yet it must be large enough to ensure safety. These questions have motivated the development of the Free Flight Transition Zone Simulation and the Controlled Airspace Model, neither of which has yet been integrated into the ASCENT testbed.

5.5.4.1 Free Flight Transition Zone Simulation

The Free Flight Transition Zone Simulation relies on two sets of simple but effective rules: one is a pilot behavior model which has been designed to assure aircraft separation and the other is a model of ATFM that influences pilot behavior in uncontrolled airspace in a manner that enforces order in the transition to controlled airspace. The rules operate the same way over Boston as they do over Nebraska, the same way at sector boundaries as in the middle of a sector, thus aircraft separation assurance need not be analyzed for a myriad of geography-specific cases as in heuristic solutions to the separation assurance problem. The separation assurance mechanism, i.e., the aforementioned pilot behavior model, used here and described below provides a model of a working separation assurance mechanism for purposes of simulating the aircraft trajectories during transition between the free flight zone and controlled airspace.

Although the separation assurance mechanism model is used in this simulation of all airspace regardless of its type of control, the mechanism has the most effect in free flight zones (since, in controlled airspace, the control ideally prevents the separation assurance mechanism from being activated). One fundamental premise of the mechanism modeled is that an effective pilot interface will have been developed that shows the pilot the set of disallowed trajectories, outside of which the pilot is free to fly. It is expected that a pilot may also have additional decision support mechanisms working in parallel to assist with the decision of which of the allowed trajectories to fly.

For the work described here, a trajectory is disallowed only if that trajectory could result in the aircraft entering another aircraft's airspace or otherwise disallowed airspace. Erzberger[10] has computed the optimal least expected distance avoidance rule for the two-aircraft case taking into account position and velocity uncertainties typical of today's air traffic information systems. The optimal course correction time he derives is well in advance of the time at which a course correction is absolutely necessary. As trajectory prediction improves (partly as a result of improved navigation) the optimal time for course correction becomes earlier still. Since airlines desire efficient operation, actual enroute flight paths under such a separation assurance mechanism would likely correct course earlier than necessary from a safety perspective, in order to avoid the inefficiency of late correction.

Our pilot behavior model can be thought of as "worst case" in that it models pilots who fail to plan ahead. The modeled pilots direct their aircraft straight for their destinations until instructed by the separation assurance mechanism that they must do otherwise. In en route

[10] Heinz Erzberger, A Presentation to the MIT Aeronautical Engineering Community, extending ideas in Russell A. Paielli and Heinz Erzberger, "Conflict Probability Estimation for Free Flight" to be published in AIAA Journal of Guidance, Control, and Dynamics, 1996.

airspace, simulation experience suggests that such pilots would cover very little distance beyond that required to follow the great-circle route; even under extremely congested conditions by today's standards, aircraft would seldom be required to fly more than one thousandth more than the great circle distance. Given wind-optimal or other desired trajectories, similar results would be anticipated.

By Erzberger's derivation, the travel times and trajectories produced in this model must be worse than optimal. Experimental results in extreme cases have shown the trajectories produced under the model to be very close in length and duration to an upper performance bound from straight trajectories. Since the performance under the model is worse than optimal yet close to the performance upper bound, it must be close to optimal. Furthermore, system performance under this model is a lower bound on the performance a real system would experience.

The separation assurance model does not provide a way to transition aircraft into a desired sequence for landing. Without additional control over the aircraft, this transition to landing could be inefficient and unsafe. In simulation, several aircraft desiring to land at the same time can interact and block the airspace leading to the airport, with neither able to land. The Controlled Airspace Model described below addresses this issue.

5.5.4.2 Controlled Airspace Model

The Controlled Airspace Model divides controlled airspace into two zones: the strictly controlled zone, in which aircraft are given precise trajectory commands by a central authority, and the transition zone, in which aircraft have more limited freedom in choosing trajectories than they have in the uncontrolled airspace. The model operates under the assumption that aircraft entering the strictly controlled zone must be temporally separated by a prespecified amount of time. The transition zone, which surrounds the strictly controlled zone, employs an ATFM-directed influence on pilot behavior that enforces this temporal separation constraint as follows.

A sequencing for the aircraft to enter the strictly controlled zone is determined by, for example, CTAS. The transition zone enforces this ordering by requiring each aircraft to stay an ATFM-specified distance away from the strictly controlled airspace as a function of the time-to-go for the aircraft to enter the strictly controlled airspace. Figure 5.5.4.2-1 shows a simple scheme in which the strictly controlled airspace is just large enough to encompass the base leg and final approach.

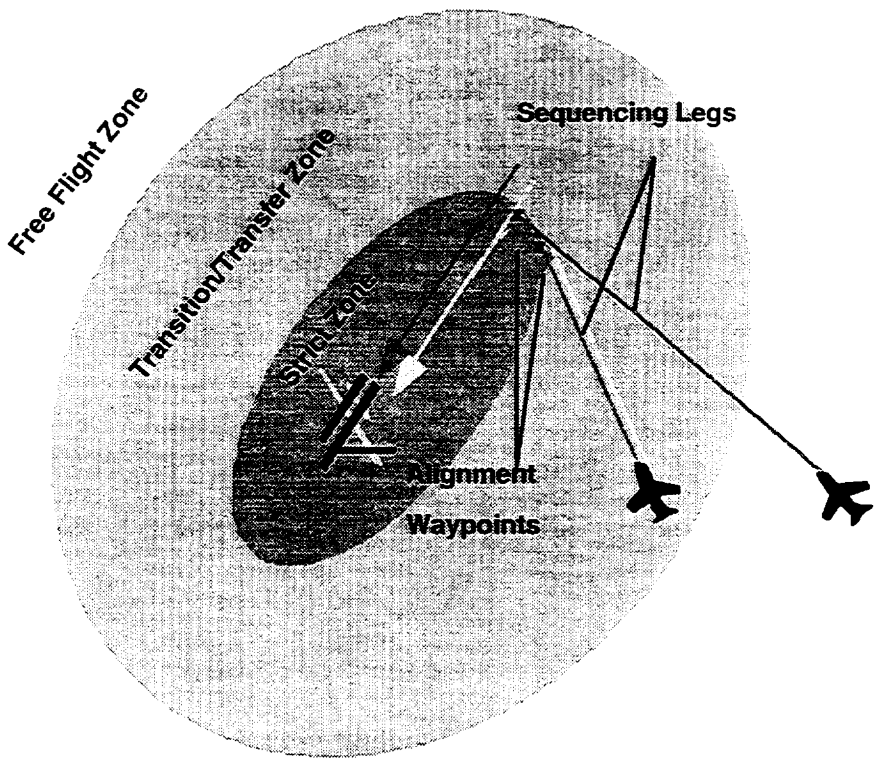

Figure 5.5.4.2-1: Simple Airport Area

Under this model, aircraft head toward their respective alignment waypoints on the boundary of the strictly controlled zone. Control of an aircraft's trajectory is transferred to the central facility at the time the aircraft reaches its alignment waypoint. The flight path before reaching the alignment waypoint within the transition zone is determined by the pilot as in the free flight zone, **except that** the pilot must also obey a keep-out constraint that prevents him from approaching the airport out of sequence (see Figure 5.5.4.2-2). The constraints are implemented in terms of travel time to the alignment waypoint rather than in terms of distance to account for aircraft of different performance. The keep-out constraint can be thought of as a collapsing sphere around the strictly controlled airspace which collapses to the strictly controlled airspace at a desired time of arrival at the alignment waypoint.

The implementation of the keep-out boundary is a constraint on each aircraft to stay a certain radius, r, from its destination, where r is a function of the time until the aircraft's slot to enter the airspace. The precise form of the function by which r is determined is a subject of this research. One candidate function (upon which the results presented here are based) is

$$r = VT\left(1 - \frac{\arg\max\{T, M\}}{2M}\right)$$

where V is the approach speed of the aircraft, T is the time-to-go until the aircraft's slot to enter airspace, and M is the maximum radius (in units of time). Thus, under this policy, the keep-out boundaries are spherical regions all of which fit within a limiting sphere around the destination. As the time approaches for an aircraft to enter the strictly controlled airspace, the sphere corresponding to that aircraft collapses at an increasing rate until it vanishes, at which point it is collapsing at the desired approach speed.

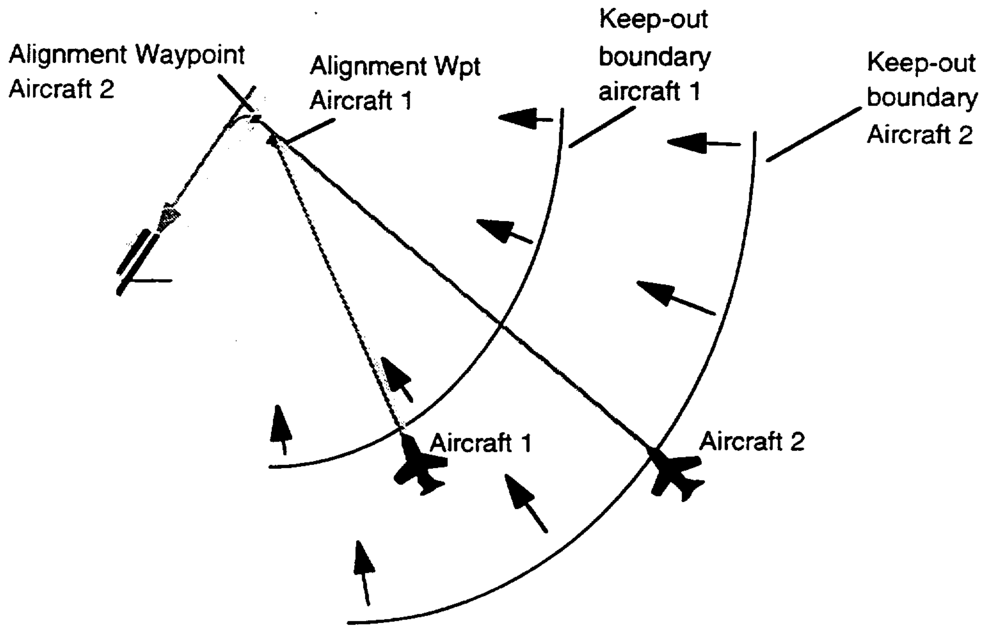

Figure 5.5.4.2-2: Keep-out Boundaries at an Airport

The initial results have been encouraging. Without this sequencing mechanism, aircraft conflicts can produce a bottleneck at the entrance to the controlled airspace, reducing the rate at which aircraft can enter the controlled airspace nearly to zero. With the mechanism in place, any practical flow rate can be achieved. However, much research remains to be done to determine the effectiveness of this method as a function of traffic density in the free flight zone, level of demand for use of the controlled airspace, and so forth. The impact of variations of the method on safety and efficiency should be studied. The best function for determining r is not known, and, indeed, for the function presented here it is not clear what the best values for M and V are under various conditions.

The simulation has demonstrated that under this model, aircraft that are constrained from entering the transfer zone wait just outside it. Under heavy traffic demands, the number of aircraft waiting outside the transfer zone can become fairly large. If the density of air traffic outside the transfer zone becomes high enough, the traffic forms a wall of sorts, preventing travel from the free flight zone into the transfer zone. Inevitably, an aircraft whose turn it is to enter the transfer zone cannot reach it immediately, and airport capacity is wasted. The solution to this problem is to increase dynamically the transfer zone radius (the parameter M above) so that waiting aircraft spread out along the circumference of the transfer zone and do not form a barrier. In simulation, we have shown that any reasonable traffic flow rate can be accommodated by using a large enough transfer zone.

5.5.4.3 More Simulation Details

In the simulation, aircraft travel in three-dimensional airspace from origin to destination. The initial headings and glide slope angles are such that each aircraft originating in the simulation in en route airspace is traveling directly toward its destination (within glide slope limits). Aircraft leaving a simulated airport do so on a predefined ascent trajectory.

The aircraft dynamics are modeled as follows. The simulation is run in one second time steps. Each aircraft is capable of a 3 degrees / sec yaw rate. It is also limited to a 0.6 degrees / sec glide slope rate, and glide slope is limited to +/- 5% grade.

The simulation records the trajectories of the aircraft and displays the distance between each pair of aircraft as a function of time (see Figure 5.5.4.3-3 for the example of aircraft waiting in airborne queue to land at an over-capacitated airport). In order to determine delays caused by heavy traffic congestion, the simulation measures the additional distance traveled (beyond the no-traffic case) and measures the number of avoidance maneuvers that are required. Additionally, tools have been developed to allow examination of the aircraft trajectories in animation and in various static displays.

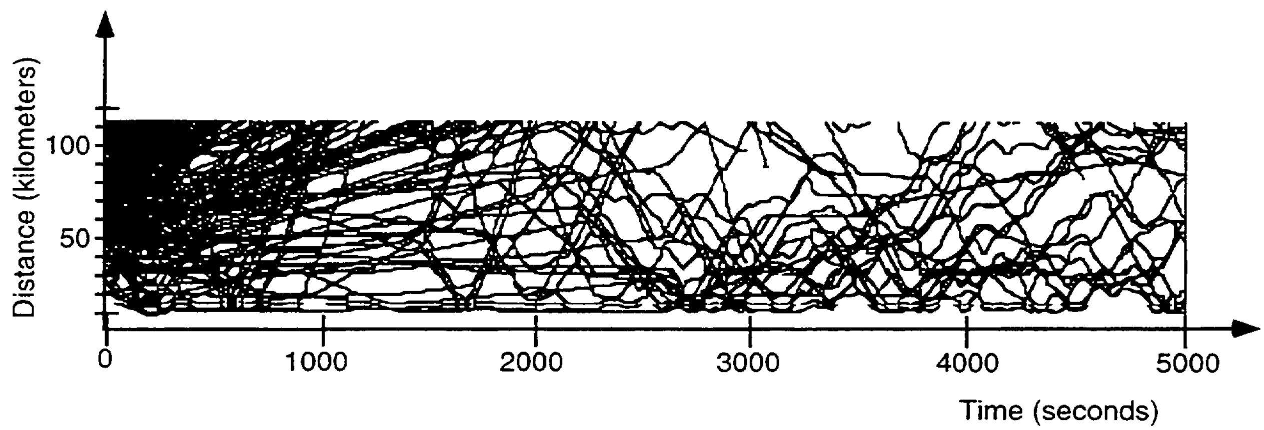

Figure 5.5.4.3-3: Aircraft Pairs Distance as a Function of Time

5.5.5 Collaborative Slot Assignment Algorithm

A heuristic-based slot assignment algorithm has been developed, based on the work of Guastalla et al.[11], that implements steps 2 and 3 of the partially decentralized scenario described in Section 5.1 and summarized here. Step 2: allocation of arrival slots at congested airports is performed on a dynamic basis, according to predicted airport capacity over the next few hours. Whenever arrival capacity at one or more airports is predicted to be scarce, available slots for arrivals at these airports are allocated among the airlines on a First-Scheduled, First-Served (FSFS) basis to ensure fairness. Step 3: each airline (or, more generally, each aircraft operator) is now free to utilize its slots in each interval in the way it deems best. Thus, each airline may schedule any one of its flights into any one of its arrival slots.

Figure 5.5.5-1 shows an example of how a typical airline applies a collaborative slot assignment algorithm[12].

---

11 Andreatta, G., L. Brunetta, G. Guastalla, "Multi-airport ground holding problem: A heuristic approach based on priority rules," Internal Report, 1 Oct. 1994.

12 The airline may first preprocess its available slots using a bank preservation algorithm like that discussed in Section 5.5.3.

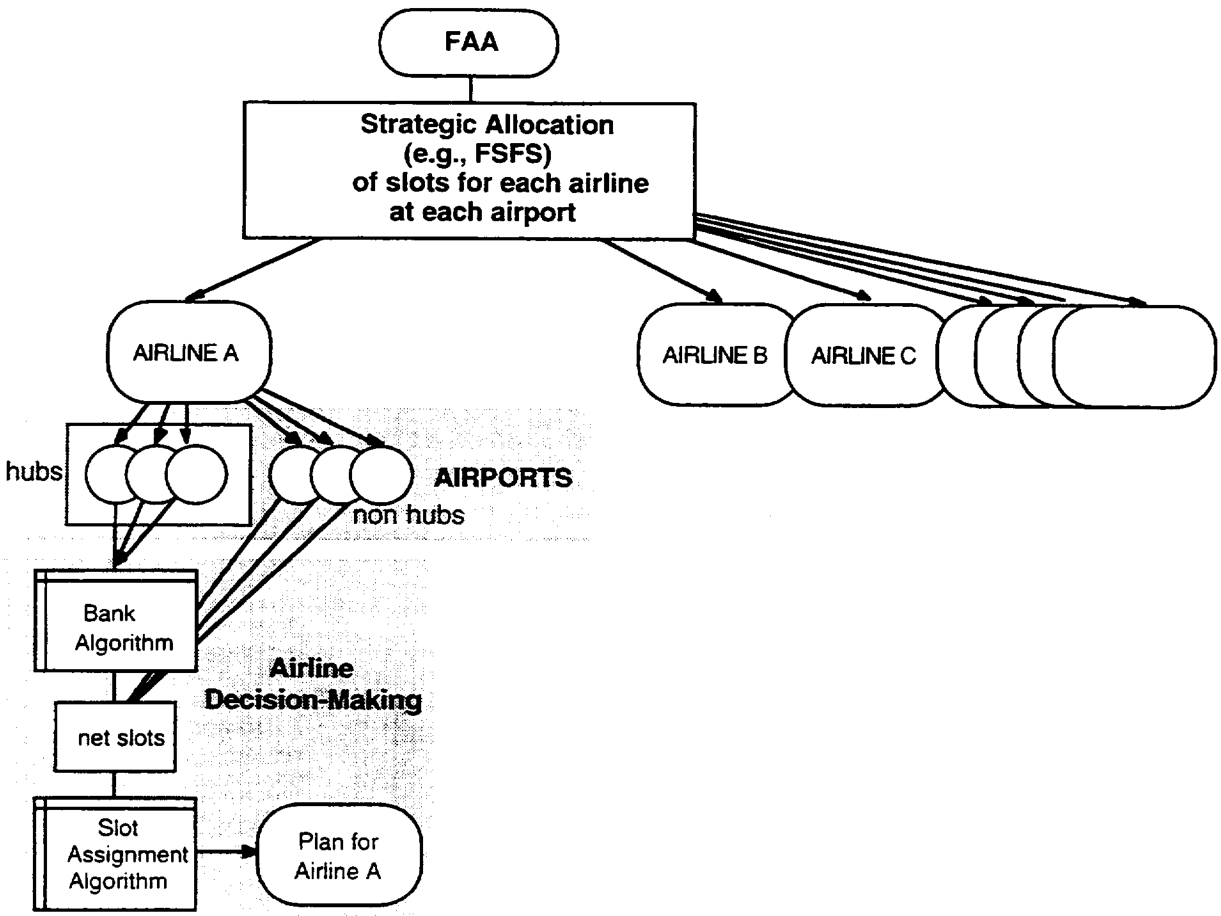

Figure 5.5.5-1 Collaborative Slot Assignment

At the heart of the airline slot assignment algorithm is the notion of a prioritization schedule; a quantized ranking of each flight that indicates its importance to the airline. The algorithm then seeks to assign any planned ground delays to the lower priority flights. For example, the results presented in Section 6 consider a priority schedule that gives a higher priority to flights employing larger passenger capacity aircraft. Many other priority schedules are possible, and in general each airline would utilize a unique prioritization that captured its own business objectives. Other example priority schedules might include assigning higher ranking to flights that have one or more of the following characteristics: flight has a connecting flight, flight has already been assigned some ground delay, flight departs from or arrives at a particular airport.

Figure 5.5.5-2 shows a flowchart of the heuristic algorithm that implements the prioritization schedule. At each airport a list of flights that would like to land at each time period is created, and ordered based on that airline's priority schedule. The algorithm then increments through the time periods, landing higher priority planes first. When capacity is exhausted at a particular time period, all unassigned flights on the list get one period of delay and are merged into the next time period's ordered list. Note that a flight's priority can dynamically change–for example, priority might contain a component that increases based on amount of delay already assigned.

The algorithm takes into consideration flight connections through the concept of critical delay. Critical delay is the amount of delay a flight can absorb and not affect the departure time of its connecting flight. The algorithm attempts to swap delays among flights so as to avoid assigning critical delays when possible.

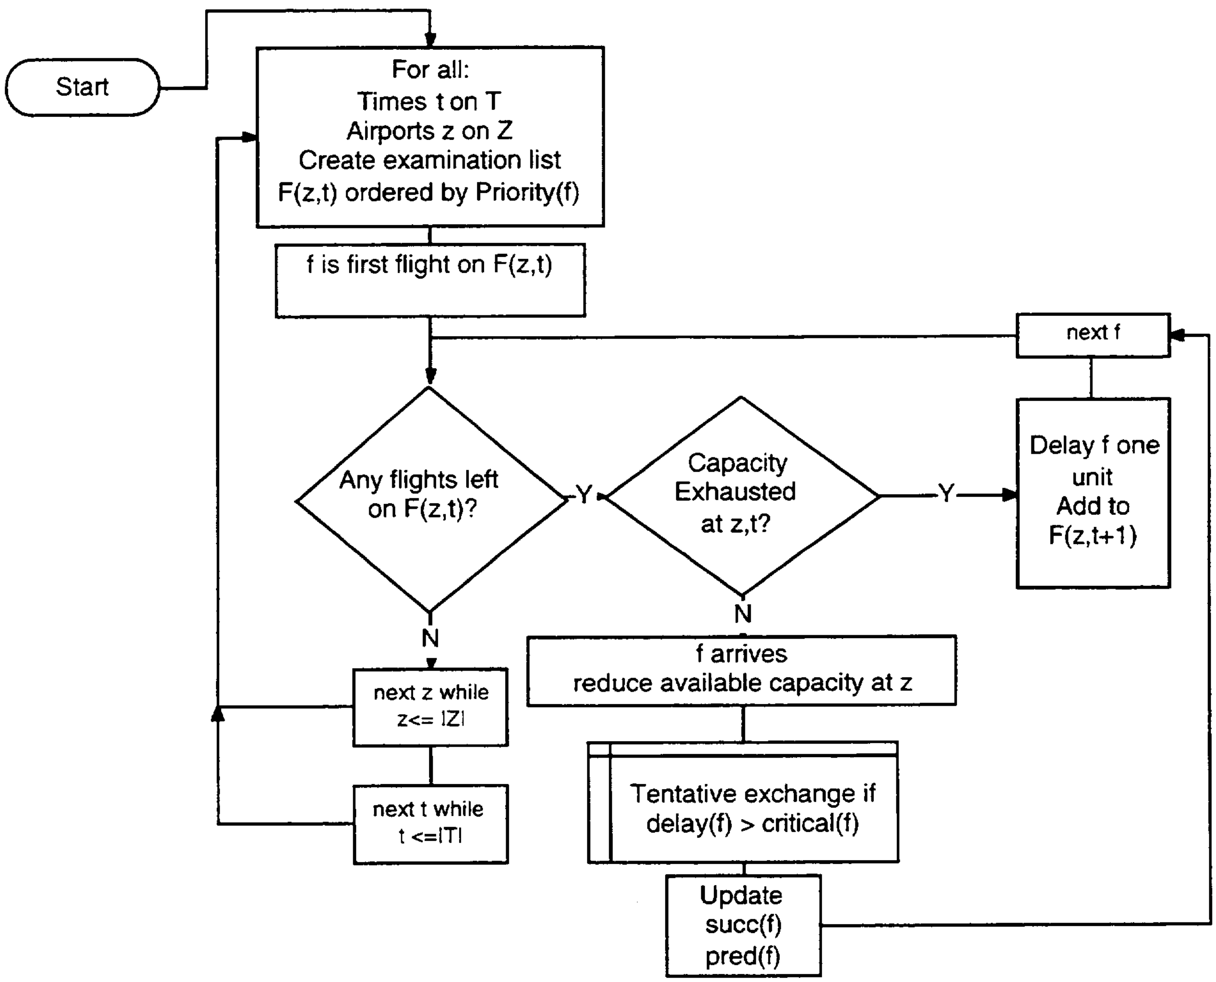

t	= a time period	T	= time periods being planned
z	= an airport	Z	= set of capacitated airports
f	= a flight	F(z,t)	= all arriving flights at time t, airport z
succ(f)	= the next flight of the aircraft that flies f	pred (f)	= the previous flight of the aircraft that flies f

Figure 5.5.5-2 Heuristic Algorithm to Achieve Airline Assignment Priorities

5.5.6 Weather and Capacity Modeling

Airport weather variables, including cloud ceiling, horizontal visibility and wind, significantly influence the capacity of an airport. Draper research on weather and airport capacity modeling started several years ago under IR&D funding and has continued under AATT funding. Robinson (1992) developed a Markov chain model suitable for single airports. Hocker (1994) and then Yu (1996) enhanced an existing validated U.S. Air Force weather model, the sawtooth wave model, a statistical weather model with the capability to generate synthetic weather observations useful in air traffic flow management analysis. The sawtooth wave model uses historical weather data as input and produces synthetic weather observations that preserve the spatial correlations of weather observations among sites in a region, temporal correlations of weather variables at each site, and cross-correlations between weather variables. The sawtooth wave model can be used to simulate weather at a set of locations in a geographical region within a radius of several hundred miles.

In his Master's thesis, Yu (1996) integrated the sawtooth wave model with several airport capacity models to analyze arrival and departure capacity patterns at DCA (Washington National), LGA (LaGuardia) and BOS (Logan). The capacity models studied were the following.

1) *Empirical Data Capacity Frontiers*—based on historical counts of arrivals and departures, developed under the FAA's Advanced Traffic Management System (ATMS) program; see Gilbo (1993).

2) *Engineered Performance Standards* (EPS)—introduced in early 1974 in an effort undertaken by the Operations Research Branch of the Executive Staff, Air Traffic Service to develop a system for measuring performance of major airports. Previous to this effort, the only indicators of an airport's performance were delay statistics maintained by airlines; see Federal Aviation Administration (1975).

3) *FAA Airfield Capacity Model*—originally developed in the 1970s by a consortium including Peat, Marwick, Mitchell, and Company and McDonnell Douglas Automation. It was later modified by the Systems Research and Development Service (SRDS) branch of the FAA. A major effort was initiated to upgrade the SRDS version to add new functions and abilities as well as incorporate the current ATC procedures. This was completed in 1981 and this is the version studied; see Swedish (1981). The model is designed to calculate the maximum throughput capacity of a runway system assuming a continuous flow of demand.

4) Runway Capacity Model—developed by the Logistics Management Institute in the NASA Terminal Area Productivity (TAP) Program; see Wingrove et al. (1995). This was done by taking airport-specific data, estimating an airport's capacity, using a queuing model to calculate aircraft delay, and subsequently calculating the cost savings to airlines by reducing delay using some or all of the TAP systems.

A weather/capacity simulation using the different capacity models was implemented and the results were analyzed.

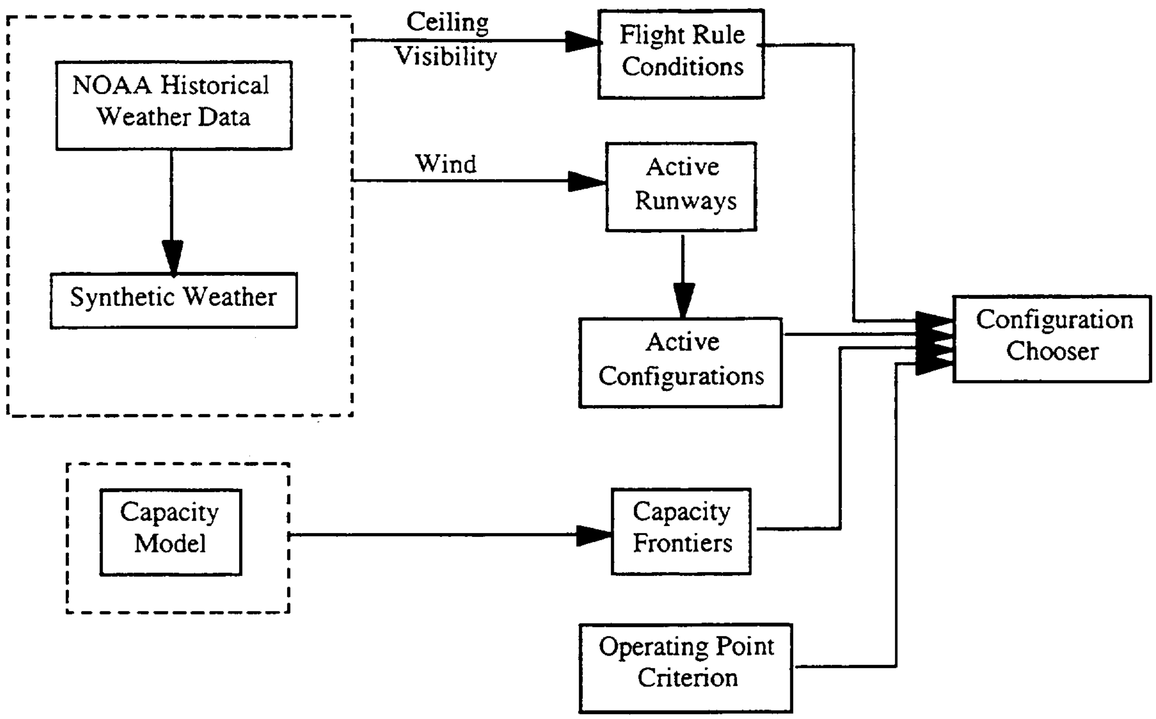

Figure 5.5.6-1: Weather/Capacity Simulation Flow Diagram

Figure 5.5.6-1 illustrates the simulation process used to obtain the results discussed in Yu (1996). The inputs to the simulation are weather observations at the airports of interest. These observations can either be historical weather data or weather synthetically generated using the sawtooth wave model. From the ceiling and visibility observations and the flight rule constraints for the airport of interest, the flight rule conditions are determined. From the wind observations, a list of active runways (runways which do not violate the maximum crosswind constraint) is generated; subsequently, a list of possible configurations using only active runways is produced. A capacity model is used to derive capacity frontiers (see Section 3.3 and Figure 4.3.3.-1) for all combinations of configurations and flight rule conditions. These frontiers, the flight rule conditions, the list of possible runway configurations, as well as an operating point criterion are inputs to a configuration chooser.

The configuration chooser outputs the configuration that best meets the operating point criterion. We used the maximum operations rate as the operating point criterion. Therefore, the configuration chooser calculates the maximum operations rate from each capacity frontier corresponding to an active configuration and current flight rule conditions. Then the configuration with the maximum of these operations rates is chosen as the best configuration.

The configuration choices and corresponding capacities produced by the simulation were analyzed in order to evaluate the accuracy of the different capacity models and identify flaws in the simulation process. To do this, three results were examined closely: configuration usage, capacity coverage, and capacity time series data. Configuration usage summarizes the percentage of time each of the different configurations for an airport is chosen. Capacity coverage illustrates how much capacity is available at an airport. Time series data is used to depict how capacity of an airport changes over time.

An example set of capacity coverage charts are shown in Figure 5.5.6-2, illustrating the seasonality of capacity coverage. During the winter, BOS operates at an operations rate of greater than 100 operations per hour only 69 percent of the time while during the summer this level of operations can be attained 80 percent of the time. See Yu (1996) for a more complete set of results.

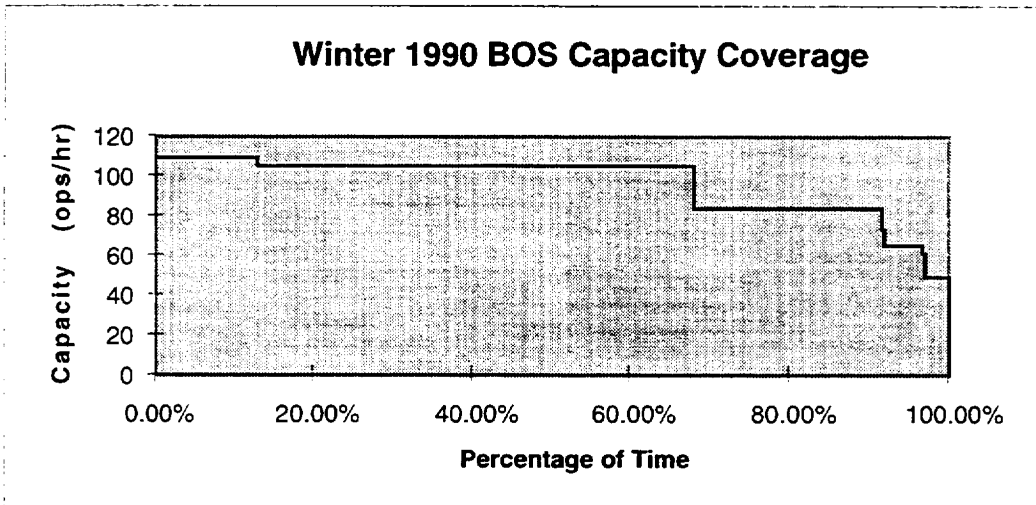

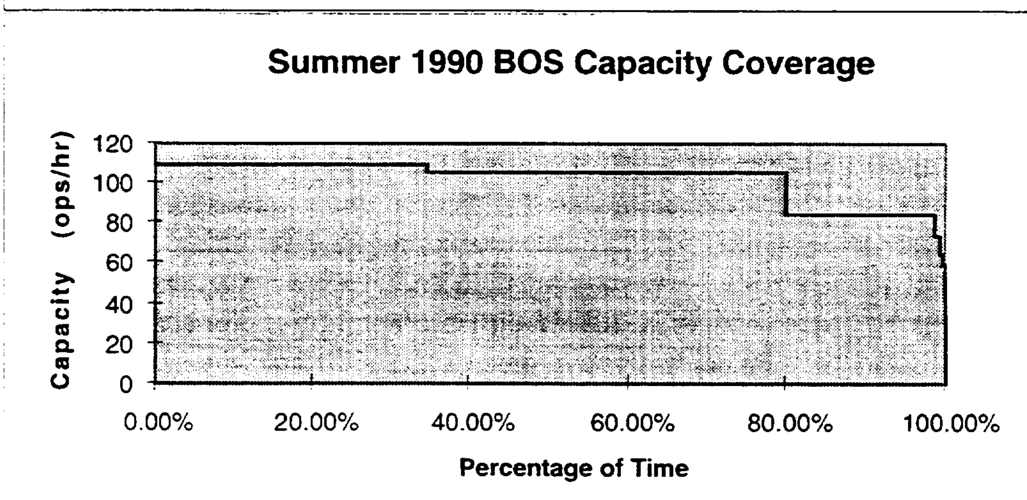

Figure 5.5.6-2: BOS Capacity Coverage Charts — Winter and Summer

5.5.7 Managed Arrival Reservoir (MAR)

The work described in this section is research in progress. The research began under AATT program funding and is continuing under Draper IR&D funding.

5.5.7.1 Overview

In air traffic management, large amounts of costs and congestion are incurred because of uncertainty relating to future landing capacity over the next several hours. Ground holding is one of the basic methods of lowering these costs. The idea is simple: it is preferable to have a flight wait on the ground at its point of origin than to have it circle the airport at its destination, unable to land. Therefore, if it is known with certainty, or at least with high probability, that a flight will be unable to land due to lack of capacity, it will be advantageous to hold the flight on the ground at its point of origin. Ground holding saves on fuel costs and increases safety margins by relieving congestion.

The FAA introduced a ground holding policy in the early 1980's. For each possibly capacitated airport, the FAA generated an estimate, or forecast, of capacity over the next few hours. The FAA then treated this forecast as a completely accurate profile of future landing capacity, and ground hold exactly enough airplanes such that if capacity materialized as planned, there would be no air holds (planes forced to wait in the air at their destination due to lack of landing capacity). This is the First-Scheduled First-Served (FSFS) ground holding policy (algorithm).

One problem with the FSFS ground holding policy is that the forecast of future capacity is treated as accurate and deterministic. In other words, once the forecast is made, the stochastic nature of future capacity is ignored.

An additional problem with the FSFS policy is that if the FAA's forecast of capacity is not equal to the *expected* values of future capacity, the FSFS policy will introduce a systematic bias in ground holding. Airlines feel that the FAA's capacity forecasts are overly conservative, corresponding more closely to worst-case scenarios than expected-case scenarios, leading to a large number of ground holds that are, from the airlines' point of view, unnecessary, and therefore lead to a large amount of wasted capacity. The Managed Arrival Reservoir (MAR) system is an attempt to deal with its difficulty.

The goal of MAR is to maintain a small reservoir of planes waiting to land at any given capacitated airport; in other words, to make sure that capacity is not wasted. In practice, this is implemented by allowing all ground held planes to depart a constant amount of time (15 minutes) earlier than they would have been allowed to depart under the original FSFS policy. This does seem to have improved the situation somewhat.

Some attempts have been made to devise alternative ground holding policies. Richetta and Odoni (1993) outlined the *static stochastic* ground holding policy. The static stochastic algorithm assumes that an accurate probability distribution of future capacity is available, and that this distribution can be represented as a manageable number of capacity *scenarios*: possible profiles of future capacity. Furthermore, the algorithm assumes that the relative costs of ground holding and air holding are known. The algorithm then formulates the ground holding problem as an integer optimization problem, using the expected cost as the objective function. The static stochastic algorithm has the advantage that it does not explicitly require a forecast of future capacity.

The static stochastic algorithm is not currently being used or considered by the FAA, for several reasons. One difficulty is that few of its assumptions currently hold in practice: an adequate capacity distribution is not easily available, and the relative costs of ground holding and air holding are unknown, and would in fact have different values to different parties—airlines would most likely want a different mix of air holding and ground holding than the FAA, for example. Another difficulty is that the algorithm has only been tested on a few idealized, simplified cases. A final difficulty is that the FAA would be extremely loath to rip out a working system and install an entirely new one; the air traffic control system is too critical to risk any interruptions in service. One of our design constraints in developing a superior MAR policy is the degree to which it differs from current operational procedure—we would like to keep the amount of change required as small as possible.

Richetta and Odoni outline a dynamic version of their algorithm, where further ground holding decisions can be made after the capacity forecast is partially realized. Due to its increased flexibility, this *dynamic stochastic* algorithm is able to achieve ground holding schedules with even lower costs than the static stochastic algorithm on simple, idealized examples. However, the dynamic stochastic algorithm has even higher barriers to being implemented in the real world than the static stochastic algorithm. In particular, one might reasonably hope that the static stochastic would perform well with an approximately correct probability distribution of future capacity, which we may hope to generate by, for example, examining historical capacity data. The dynamic stochastic algorithm, on the other hand, very much requires an exact distribution. This is because a decision at time t in the dynamic algorithm is made by conditioning on which capacity scenarios match the realized capacity up to time t; if no scenarios match, the algorithm fails. It is extremely difficult to obtain an exact probability distribution of future capacity; even if it were done, it is likely that this distribution would consist of a unmanageable number of cases.

The goal of this research is to evaluate the current implementation of MAR, and attempt to determine a better MAR policy by varying the amount by which ground holds are reduced, possibly allowing differing ground holds as a function of airport, time of year, and initial weather conditions. We will require that the reduction in ground holds be the same for all ground-held aircraft: given two airplanes being ground held for one hour each under the FSFS policy, we will not allow ourselves to ground hold one plane for fifteen minutes and the other for a half hour.

5.5.7.2 Methodology

If we are unable to model in any way the future distribution of weather from a given point in time, we are essentially stuck. We can do no better than assume that our *forecast* is correct and ignore the stochastic nature of the problem. Furthermore, without being able to model the stochastic distribution of capacity, we are unable even to *evaluate* the expected cost of any particular policy. For this reason, we begin by focusing on capacity distributions.

The first-order determinant of airport capacity is weather, including visibility, wind direction, speed, and cloud ceiling, but other factors, including the experience of the pilot and the air traffic controllers on duty, or the current mix of small and large aircraft attempting to land, also affect capacity. Nevertheless, for the combined reason that weather is the most important determinant of capacity and that weather data is the only data we actually have available, we chose to model weather as the sole determinant of airport capacity.

We obtained historical data on cloud ceiling, visibility, wind speed and direction from a commercially available National Oceanic and Atmospheric Administration (NOAA) CD-ROM. The NOAA CD-ROM contains hourly data from 1961-1990, but there are many missing observations. Since the missing data points follow the pattern that for many years in many locations, observations were only being taken every three hours instead of every hour, we assumed there was no systematic bias to the missing observations, and discarded all incomplete days.

Next, we used the cloud ceiling and visibility data to determine a flight rule for each observation. There are four different flight rules: VFR1, VFR2, IFR1, IFR2, listed in decreasing order of visibility and cloud ceiling requirements.[13]

Next, we used a model developed by the FAA that maps the flight rule and the wind speed and direction at a given airport into an arrival capacity value. Strictly speaking, the FAA model maps weather conditions into a number of different possible configurations, where each configuration consists of arrival and departure capacities. We chose a point on this frontier that optimized arrival capacities, on the assumption that arrival congestion is more significant than departure congestion.

We segregated data by month and used as our empirical distribution all days at a given airport in a given month, with a given flight rule and capacity at the start time. For example, if, at 6 a.m. on July 7, we look outside in San Francisco and the current flight conditions are IFR2, with a landing capacity of 30, we would use as our capacity distribution all observed days in July in San Francisco that had IFR2 flight conditions and a landing capacity of 30 at 6 a.m.

If we are going to test the performance of any algorithms, we need to know the costs of air and ground holding. Unfortunately, as was stated above, these are not known. We can, however, try out a range of different air and ground holding costs, and hope to find an optimal policy as a *function* of the ground and air hold costs. We further note that the absolute values of the costs are not important, and that we may therefore achieve our aim merely by varying the *ratio* of air to ground holding costs. For this particular set of experiments, we tried ratios of 1.5, 2.0, 2.5, and 3.0.

Additionally, in order to apply the FSFS algorithm and its MAR variant (but not the static stochastic algorithm), we need a *forecast* of future capacity. We would like to have a forecast that accurately matches what the FAA might predict given some set of initial conditions, but we have no data on FAA forecasting. We approach this issue by using three different forecasts. Two of these forecasts are an *optimistic* forecast and a *pessimistic* forecast. For each case, we look at all days matching the initial conditions, i.e., our empirical distribution for a given case. For the optimistic forecast, we choose the highest observed capacity at each time, and for the pessimistic forecast we choose the lowest observed capacity at each time. For example, if the empirical distribution for San Francisco in July, with flight rule VFR1 and starting capacity 38 consisted of three profiles, the first being 9, 7, 9, 8, 5, the second being 10, 6, 6, 9, 6, and the third being 8, 8, 8, 7, 10, the optimistic forecast would be 10, 8, 9, 9, 10, and the pessimistic forecast would be 8, 6, 6, 7, 5. The final forecast, which we call a *one-third* forecast, chooses, for each case, a value that is the n/3rd order statistic of the capacities at each period, where n is the number of observations in the capacity scenario.

In all cases, we discretized time into fifteen minute intervals. Our capacity forecasts were hourly, so we divided the hourly capacity into four, "backloading" any remainder. For instance,

---

[13] VFR stands for *visual* flight rules, while IFR stand for *instrument* flight rules.

if the capacity for a given hour were 38, the capacity for the four fifteen-minute periods in that hour would be 9, 9, 10, and 10, in that order.

For demand data, we used the OAG data for Jan 13 and July 3, 1993.

5.5.7.3 Results

We ran the FSFS algorithm on each possible combination of forecasts (optimistic, pessimistic, or one-third), airports (Boston, Dallas-Ft. Worth, or San Francisco), dates (Jan 13 or July 3), ratio of air holding to ground holding costs (1.5, 2.0, 2.5, or 3.0), and starting conditions (flight rule and capacity at that airport). In all cases, we assumed that we had made the 6 a.m. observation, and were trying to assign ground-holds for 7 a.m. to 2 p.m. In each case, we tried all possible values of the MAR parameter, from 0, which is simply the FSFS policy, up to a number high enough that no planes were being ground-held. Rather than present the data in its entirety, we summarize the data, present a few typical examples, and attempt an interpretation.

The optimistic forecast was very easy to analyze. Considering a "case" to be a specific airport, date, and initial capacity, there were a total of 64 cases. The optimal MAR value under the optimistic forecast was identically zero across all 64 cases. Therefore, we feel comfortable suggesting that if the forecasting methodology were essentially generating optimistic forecasts, using a MAR value of 0 (i.e., simply using the original FSFS algorithm) would be preferable to the current MAR value of 1. Of course, we do not believe that the FAA is generating optimistic forecasts; one of the primary reasons for the implementation of MAR in the first place was that the FAA's forecasts were seen as too conservative, leading to wasted capacity.

The one-third forecast proved more interesting. Here again we find a large number of zero values: 53 of the 64 cases had an optimal MAR of 0 regardless of cost ratio. Nine of the eleven cases had a non-zero MAR value only when the cost ratio was 1.5, and in these cases the optimal MAR was 1. We summarize these cases here:

Airport	Flight Rule	Initial Capacity	Month	MAR=0 Cost	MAR=1 Cost	% Savings
BOS	IFR2	34	JAN	63,840	61,382	4.0
BOS	VFR1	56	JAN	3,930	3,806	3.3
DFW	IFR1	60	JAN	153,575	150,450	2.1
DFW	IFR2	50	JAN	178,787	175,765	1.7
DFW	VFR1	95	JUL	30,813	28,300	8.9
DFW	VFR2	66	JUL	66,352	65,805	0.9
DFW	IFR1	60	JUL	66,785	63,128	5.8
SFO	IFR2	34	JAN	43,115	41,201	4.6

For the two cases that had a non-zero MAR value for cost ratios higher than 1.5, we include the actual output, both to analyze the specific cases and to familiarize ourselves with the form of the underlying data. The first case is from output file BOS.JAN.mar13:

dist.V1.42,1.5:

Mar = 0	Ground Cost = 8100	Air Cost = 1387	Total = 9487
Mar = 1	Ground Cost = 0	Air Cost = 7187	Total = 7187

dist.V1.42,2:

Mar = 0	Ground Cost = 8100	Air Cost = 1849	Total = 9949
Mar = 1	Ground Cost = 0	Air Cost = 9583	Total = 9583

dist.V1.42,2.5:

Mar = 0	Ground Cost = 8100	Air Cost = 2312	Total = 10412
Mar = 1	Ground Cost = 0	Air Cost = 11979	Total = 11979

dist.V1.42,3:

Mar = 0	Ground Cost = 8100	Air Cost = 2775	Total = 10875
Mar = 1	Ground Cost = 0	Air Cost = 14375	Total = 14375

We see that in this case, 0 and 1 were the only possible MAR values --- the FSFS algorithm does not ground hold any planes for longer than 15 minutes in this case. Additionally, although a MAR value of 1 is optimal for cost ratios of 1.5 and 2, the advantage of a MAR of 1 is 32% if the ratio is 1.5, but only 3.8% if the ratio is 2.0.

The other case of interest is Dallas in January, with an initial flight rule of VFR1 and an initial capacity of 90:

dist.V2.90,1.5:

Mar = 0	Ground Cost = 71200	Air Cost = 51200	Total = 122400
Mar = 1	Ground Cost = 29300	Air Cost = 70714	Total = 100014
Mar = 2	Ground Cost = 7500	Air Cost = 90642	Total = 98142
Mar = 3	Ground Cost = 800	Air Cost = 100192	Total = 100992
Mar = 4	Ground Cost = 0	Air Cost = 101392	Total = 101392

dist.V2.90,2:

Mar = 0	Ground Cost = 71200	Air Cost = 68266	Total = 139466
Mar = 1	Ground Cost = 29300	Air Cost = 94285	Total = 123585
Mar = 2	Ground Cost = 7500	Air Cost = 120857	Total = 128357
Mar = 3	Ground Cost = 800	Air Cost = 133590	Total = 134390
Mar = 4	Ground Cost = 0	Air Cost = 135190	Total = 135190

dist.V2.90,2.5:

Mar = 0	Ground Cost = 71200	Air Cost = 85333	Total = 156533
Mar = 1	Ground Cost = 29300	Air Cost = 117857	Total = 147157
Mar = 2	Ground Cost = 7500	Air Cost = 151071	Total = 158571
Mar = 3	Ground Cost = 800	Air Cost = 166988	Total = 167788
Mar = 4	Ground Cost = 0	Air Cost = 168988	Total = 168988

dist.V2.90,3:

Mar = 0	Ground Cost = 71200	Air Cost = 102400	Total = 173600
Mar = 1	Ground Cost = 29300	Air Cost = 141428	Total = 170728
Mar = 2	Ground Cost = 7500	Air Cost = 181285	Total = 188785
Mar = 3	Ground Cost = 800	Air Cost = 200385	Total = 201185
Mar = 4	Ground Cost = 0	Air Cost = 202785	Total = 202785

Here the value of MAR is non-zero for all the cost ratios we tested. The optimal MAR value is 1 for cost ratios of 2.0, 2.5, and 3.0, and is 2 for a cost ratio of 1.5; this is the *only* example we found where the one-third forecast led to an optimal MAR value greater than one.

Overall, we found that under the one-third forecasting policy, the optimal MAR value was nearly always zero, and in nearly all of the cases when it was not zero, the savings by using the optimal MAR rather than a MAR value of zero were fairly small. Additionally, the cases where the optimal value was non-zero did not appear to be systematic in any way: they came from all possible flight rules, and within each flight rule, the initial capacities were not always at the bottom or the top of the range (for example, for Dallas in July, 95 is the highest possible starting capacity for VFR1, whereas 66 is the median starting capacity for VFR2). Therefore, as with the optimistic forecasts, we suggest that if the forecasting method used in practice corresponds closely to one-third forecasts, a MAR value of 0 (corresponding to the use of the FSFS algorithm) is likely to be superior to a MAR value of 1.

We now turn to an analysis of the pessimistic forecast. This is the most difficult case, because the data varied so widely. Out of the 64 cases, there were only six cases where the MAR value was identically zero. We list them here:

Airport	Month	Flight Rule	Initial Capacity
BOS	JAN	IFR2	31
BOS	JUL	VFR1	42
DFW	JUL	VFR2	30
SFO	JUL	IFR1	37
SFO	JUL	IFR1	43
SFO	JUL	IFR2	37

It is difficult to make any useful conclusions from the above table. Five of the six cases occur in July, but this hardly seems statistically significant. Possibly more interesting is that five of the six cases (all except the SFO IFR2 case) involve starting capacities at the bottoms of their respected ranges (for example, 30 is the lowest possible starting VFR2 capacity for DFW in JULY). This would be far more interesting if it were coupled with a trend for lower starting capacities to lead to lower MAR values in general. However, this does not seem to be the case.

We now turn to another interesting set of cases, those cases where the MAR value was identically equal to the maximum possible value across the cost ratios tested. The implication is that in these cases, the optimal MAR value is such that there are no ground holds; allowing all planes to take off as scheduled in this manner is known as the *passive* algorithm in the literature. There were five cases that fit this pattern:

Airport	Month	Flight Rule	Initial Capacity	MAR Value
DFW	JAN	VFR1	80	30
DFW	JUL	VFR1	95	28
SFO	JAN	VFR1	52	4
SFO	JAN	VFR1	56	4
SFO	JUL	VFR1	56	2

We immediately note that all five cases occurred under VFR1 conditions. Additionally, four out of the five cases (all but the DFW JAN case) were at the top of their respective ranges. In the fifth case, for Dallas in January with a starting flight rule of VFR1 and capacity of 95, the optimal MAR was the maximal possible, 11, for cost ratios of 1.5 and 2.0, dropping to 9 for a cost ratio of 2.5, and 8 for a cost ratio of 3. Furthermore, in this case, the difference in costs between using a MAR value of 11 and using the optimal MAR value is very small. All this would seem to suggest a policy guideline that in San Francisco or Dallas, under the pessimistic forecast, under VFR1 conditions at or near the top of the range, a policy of no ground holds whatsoever is essentially optimal. We are unclear why such a policy might work in San Francisco and Dallas but not in Boston. It also seems that this result is probably not very useful, since the pessimistic forecast seems especially unlikely to be generated on maximal capacity VFR1 days.

The rest of the data does not prove terribly fruitful. Of the remaining 53 cases, the MAR value is constant across the range of cost ratios tested in only one case: in San Francisco in July, with an IFR1 flight rule and starting capacity of 34, the MAR value is always 1. In the other 52 cases, the MAR value decreases over the range of cost ratios tested. It decreases at varying rates in varying cases. For instance, in Dallas in January, with a starting flight rule of IFR2 and starting capacity of 60, the optimal MAR value is 7 with a cost ratio of 1.5, but 0 for all other cost ratios tested. On the other hand, in Boston in July, under VFR1 conditions and a starting capacity of 56, the optimal MAR value was 3 for cost ratios of 1.5 and 2, and 2 for cost ratios of 2.5 and 3. These differences do not seem to be systematic in any way, and therefore do not lead directly to policy suggestions.

Finally, we make note of several checks we can perform that indicate the correctness of our algorithm and give us some insight into the problem. The first is that for any given case, the expected costs for the highest value of MAR with respect to all three forecasts should be equal; all correspond to a policy of no ground holds whatsoever, and should produce identical schedules. As a second check, we should expect that for a given case, the optimal value of MAR should never increase as we increase the ratio of air to ground holding costs.[14] This is in fact the case. Finally, we note that for the *pessimistic* forecast only, a MAR value of 0 should have expected (and actual) air costs of 0: since we estimated the capacity absolutely conservatively, and then held planes absolutely conservatively, there is no way we can experience any air delays in this case, under the assumptions of our model. This check also holds true on our data.

5.5.7.4 Conclusions and Future Directions

We built an empirical model of future capacity, and used this to evaluate the performance of the FAA's FSFS algorithm for various values of the MAR parameter under the current implementation of MAR. We performed this evaluation for three types of forecast: an optimistic, best-case forecast, a pessimistic, worst-case forecast, and a "one-third" forecast, which essentially consists of moving one-third of the way from the pessimistic towards the optimistic forecast. We saw that for the optimistic forecast, the optimal MAR value was 0. This makes sense intuitively: the optimistic algorithm will produce very few ground holds, so reducing these ground holds further is unlikely to prove fruitful. The pessimistic forecast was more difficult to analyze. No single value of MAR seemed optimal over a wide range of cases, though large positive values of MAR were often indicated. The one-third case was the most interesting. At a guess, the one-third case seems likely to most accurately model actual forecasting technology—conservative but not incredibly so. In the one-third case, we presented strong evidence for setting the value of the MAR parameter to 0. This indicates that under the assumptions of our model, including our particular choice of objective function (expected costs of air and ground holds), the current implementation of MAR may be doing more harm than good.

A large number of future directions are suggested by this research. The current study only examined expected values. It would be interesting and not terribly difficult to generate variances for each case. Also, given our current objective function, instead of merely examining a few specific ratios of air to ground costs, we could employ a form of parametric programming to completely determine the functional dependence of the optimal MAR value as a function of the ratio of air to ground cost. This would allow us to characterize more accurately and more completely the sensitivity of the MAR values to the cost ratio in various cases. Of course, some attempts to determine at least good ballpark estimates of the actual values of this ratio to varying parties would also be in order. Finally, it is not clear that our current objective function,

---

[14] Proof: Let the optimal MAR value be x for some cost ratio r. (If there are multiple optimal values, let x be the *largest* of them; our proof then shows that no larger value of MAR ever becomes optimal as r is increased). The expected cost for MAR value x is denoted by $g_x + r^*a_x$, where g_x denotes the expected number of ground holds, and a_x denotes the expect number of airholds, given MAR value x. Consider some $y > x$. The ground holding costs are a monotonically decreasing function of the MAR value, so $g_y < g_x$. Although the air holding costs are not necessarily a monotonically *increasing* function of the MAR value, the fact that x is an optimal value implies that $a_x < a_y$ because $g_x + r^* a_x <= g_y + r^* a_y$ must hold, which implies that $g_x - g_y <= r^*(a_y - a_x)$, which implies that $a_y - a_x >= (g_x - g_y)/r$, the desired result. Now we can easily see that if we increase the value of r, the final relation above must still hold, completing the proof.

Furthermore, as long as no value of MAR has expected *air* costs less than or equal to those for MAR=0, then, for all high enough ratios of air to ground costs, the expected MAR value will be 0: if $a_0 < a_y$, then $g_0 + r^* a_0 < g_y + r^* a_y$ whenever $r > (g_0 - g)/)/(a_y - a_0)$.

which only includes terms relating to expected air and ground holding costs, is appropriate. The goal of MAR is to minimize *wasted* landing capacity, so perhaps we should include such a term explicitly in the objective function. This can be easily done for the FSFS algorithm and the MAR adjustments. Intuitively, we would expect that if the objective function had a high enough coefficient for this term, not wasting landing capacity would become the most important factor, and a policy of no ground holding would therefore become optimal. The practical implications of such an analysis would of course depend critically on coming up with reasonable estimates of the relative costs of air holding, ground holding, and wasted capacity, a difficult task.

Additionally, our model is crude, and can probably be refined. For example, some of the empirical distributions we examined were based on an extremely small number of actual data points; in several cases, there was only actually a single day matching the given initial conditions, so the distribution was deterministic, and matched precisely the forecast. Although the cases that were based on a very small number of days did not match up in any systematic way with the results, we might hope to achieve more uniform results by using more data. We might consider using a sliding window of months: for example, if we were interested in calculating an empirical distribution for weather in January, we might use historical data from December, January, and February. Additionally, it is possible that we've overconditionalized, giving us too many cases with too few points each; some sort of clustering scheme may alleviate this problem.

6. CASE STUDIES

There are three major components of the work reported here. The first is the development of alternative Air Traffic Flow Management (ATFM) concepts ranging from centralized to decentralized. The second is the development of metrics and tools to be employed in the analysis of alternative ATFM concepts that are in this range. The third is the application of these tools and metrics in the analysis of three ATFM operational concepts: passive, current, and collaborative. *Passive* refers to no ATFM and is used as a baseline; *current* is the currently employed approach to ATFM; and *collaborative* refers to a hypothesized ATFM concept that is more decentralized than current ATFM and is a concept in which the FAA allocates a set of arrival slots for each airline and each airline individually determines the assignment of each of its particular aircraft to each slot in its allocated set of slots. The collaborative scenario is one that realistically could be implemented in the medium term; many of its elements are ready to be put in place today. Comparative results for each of these three concepts were generated using January 13, 1993 OAG data as a representative schedule. Five airports were chosen as the capacitated network: BOS, EWR, LGA, DCA, and PIT. All flights arriving at, or departing from, any of the capacitated airports constituted the subset of the OAG schedule used for analysis, a total of 4925 flights. Two classes of capacity[15] scenarios are used. The first is a case for which there is VMC throughout the system, called the "blue sky" scenario. The second represents the case of a weather front moving up the east coast, causing IMC for several hours (with uncertain start times) at key airports; this scenario is called the "weather front" scenario, and is detailed in Table 6-1. Cases with and without en route free flight (User Preferred

[15] EPS (Section 5.5.6) capacity values are used

Routing) and CTAS are modeled. Appropriate metrics (e.g., delay, tardiness) for each of the operational and capacity scenarios have been calculated and are presented.

Probability	Affected Airports	Time that 3 Hour IMC Period Begins
.25	PIT & DCA	Noon
	EWR & LGA	2 PM
	BOS	4 PM
.5	PIT & DCA	1 PM
	EWR & LGA	3 PM
	BOS	5 PM
.25	PIT & DCA	2 PM
	EWR & LGA	4 PM
	BOS	6 PM

Table 6-1: Description of Weather Front Scenario.

Passive results provide a baseline for the measurement of the effects of new technologies and policies on the Air Traffic Management system. Delay and tardiness statistics in current and future ATFM can be viewed as additional delays and tardiness relative to that in a passive system. The delays and tardiness in a passive system are a result of airline schedules being flown against airport capacity. It is assumed that only ATC takes place, no ATFM. This is an interesting case, since examination of system demand (OAG schedule) versus supply (airport capacity) demonstrates that delays are inevitable even without assigning delays through ATFM.

In order to analyze the benefit of the collaborative slot allocation algorithm, as discussed in Section 6.3, the flights from the January 13, 1993 OAG data are categorized by their passenger capacity. Table 6-2 displays the number of planes in each size category for all the airlines in the dataset, and for two specific airlines that are designated in this report as Airline XX and Airline YY.

Airplane Capacity	All Airlines	Airline XX	Airline YY
< 50	1360	346	472
50 –99	528	0	239
100 – 149	1366	28	927
150 – 199	1172	201	326
200 – 249	262	84	27
250 – 299	92	42	15
≥ 300	145	16	0
TOTALS	**4925**	**717**	**2006**

Table 6-2: Distribution of Aircraft by Passenger Capacity.

The collaborative scenario investigated here can be viewed as a possible next step in the transition of the ATFM system toward increased decentralized collaborative decision making. The kinds of analyses the tools we have developed under this effort make possible for this and other anticipated candidate approaches will lead to more systematic and rigorous approaches to the evaluation of the risks and benefits of proposed approaches. The ability to perform comparative analyses is considered to be the most significant contribution of our research and development effort.

The metrics used in this section are defined in Figure 6-1. Delay is unambiguously defined, but was defined independent of FAA delay definition. Defining delay can be tricky, depending upon the point of view. For example, suppose an aircraft arrives nears its destination airport well before its scheduled time, then is held locally in the air, but arrives at the gate on time. For the passenger who is not connecting to another flight, there was a delay (assuming he or she finds out about the air hold). A passenger connecting to another flight might not feel that there was delay. The airline might or might not regard it as a delay, depending upon the fate of the aircraft.

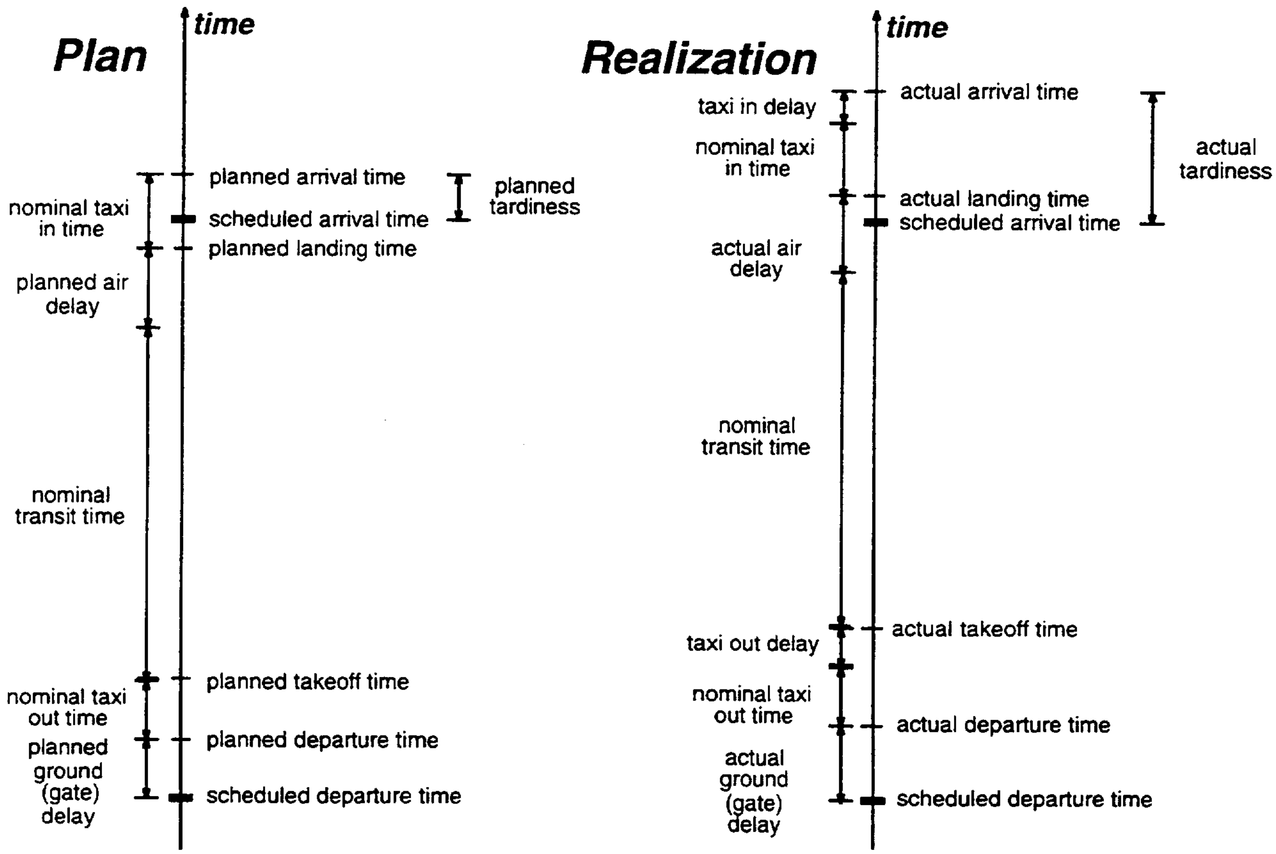

Figure 6-1: Tardiness and Delay—Planned and Realized

6.1 Effect of Free Flight and CTAS

This section addresses system level assessment of the potential benefit to be gained from free flight and CTAS. For the results presented here, free flight (user preferred routing) is modeled in ASCENT as a user-specified decrease in transit times, and CTAS is modeled as a user-specified increase in airport capacity. The specific values used are 5% for transit time reduction and 15% for capacity increase respectively. Of course, both of these percentages could be varied to determine system performance sensitivity to reduced flight times and increased airport capacities. Current ATFM was run against the blue sky scenario and the weather front scenario. Two metrics are presented in each case: the average actual tardiness of all flights and the percentage of flights that were tardy.

6.1.1 Blue Sky

Figure 6.1.1-1 shows that the introduction of CTAS and UPR reduced the average tardiness by at least 50% in all passenger capacity categories (e.g. average tardiness for 50 passenger flights dropped from 3 minutes to 1 minute, for a 66% reduction); total tardiness was reduced by 56% in this case (i.e., from 172,980 minutes to 75300 minutes). In Figure 6.1.1-2, the percentage of flights actually tardy was reduced in absolute percentage by at least 10%, for the three smallest aircraft capacity categories (two thirds of all flights belong to these three categories). For example, the percent flights tardy dropped from 29% to 16% for 50 passenger flights.

A drop off in tardiness with larger capacity aircraft is observed. Under blue sky conditions, secondary schedule characteristics influence the character of the results. Figure 6.1.1-3 displays the number of flights that have connections for each passenger capacity category. Tardiness, although small in magnitude, is propagated through the network. This propagation exists under the weather front situation as well, but arrival capacity shortages dominate the average tardiness measures in those cases.

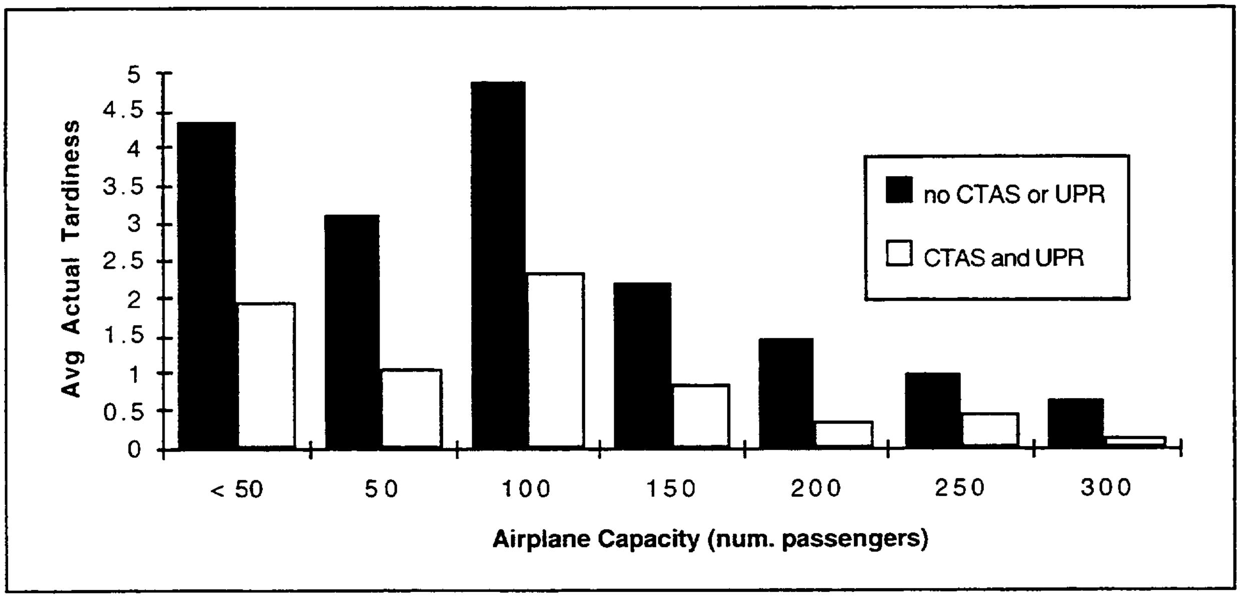

Figure 6.1.1-1: Benefit of CTAS and UPR in improving average tardiness in blue sky

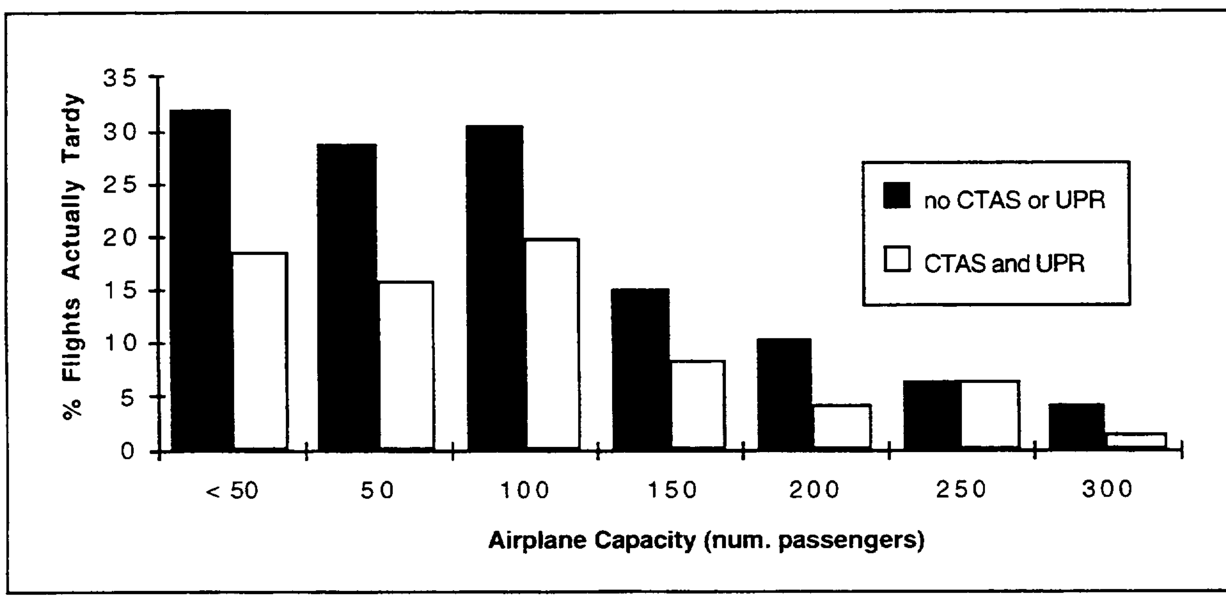

Figure 6.1.1-2: Benefit of CTAS and UPR in improving % of tardy flights in blue sky

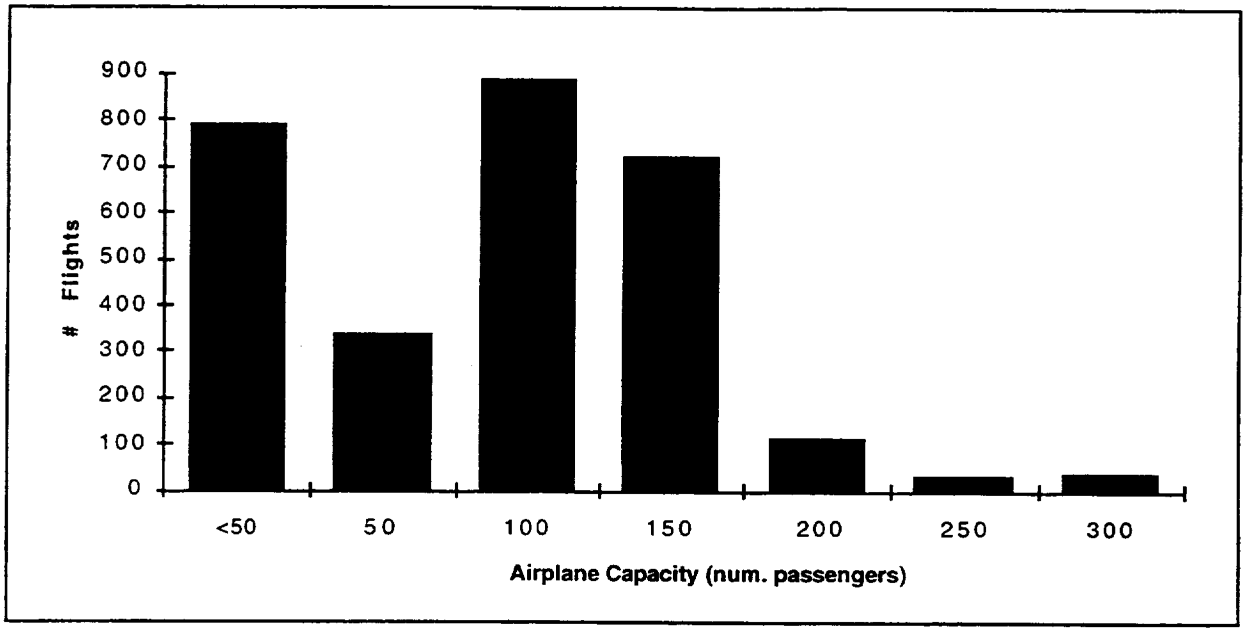

Figure 6.1.1-3: Number of flights with connections

6.1.2 Weather Front

In Figure 6.1.2-1 CTAS and UPR are shown to reduce the average tardiness by at least 32% in all aircraft capacity categories; not as big a drop as in the blue sky case. In Figure 6.1.2-2, the percentage of flights actually tardy was reduced by at least 8%, in absolute percentage, in all aircraft capacity categories. This highlights the fact, once again, that a major system performance driver is airport capacity and when that capacity is reduced, the benefits of other system improvements (in this case CTAS and UPR) are also reduced.

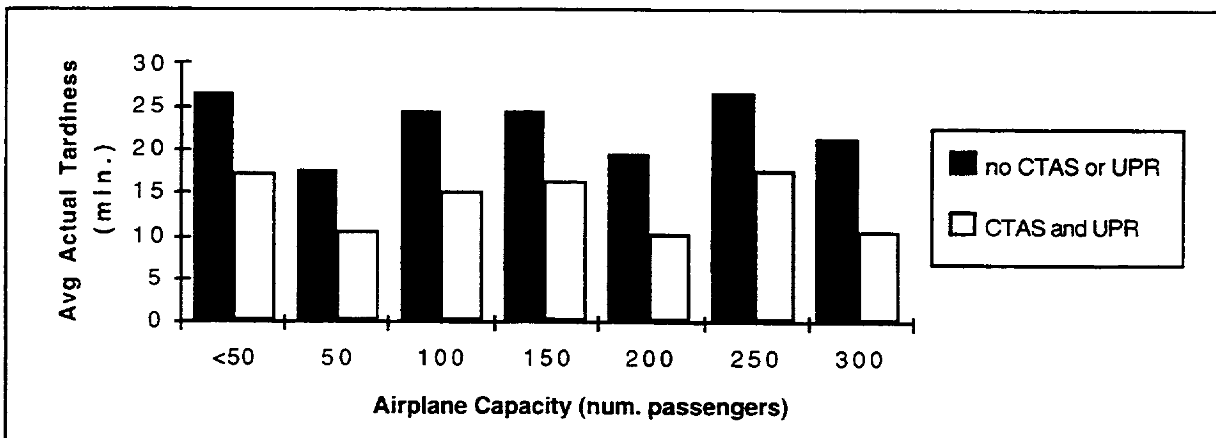

Figure 6.1.2-1: Benefit of CTAS and UPR in improving average tardiness in a weather front

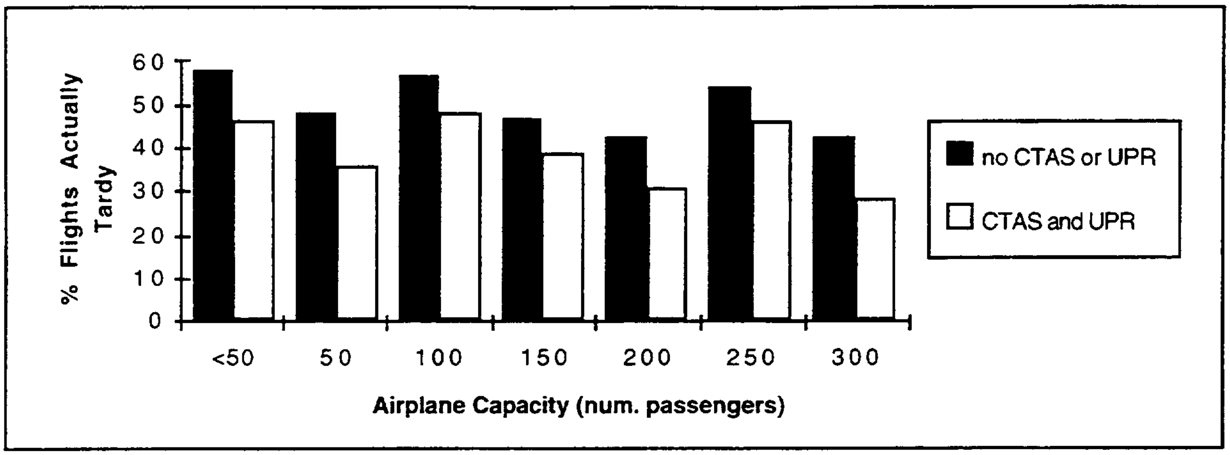

Figure 6.1.2-2: Benefit of CTAS and UPR in reducing % of tardy flights in a weather front

6.2 EFFECT OF ATFM

This section assesses the benefit effects of using ATFM. CSA (Collaborative Slot Assignment) is the partially decentralized ATFM algorithm described in Section 5.5.5.

6.2.1 Blue Sky

As explained in Section 6.1.1., the apparent preference for the larger aircraft for no ATFM and current ATFM are a result of network effects from connections. At first examination, it appears from Figures 6.2.1-1 and 6.2.1-2 that ATFM does not provide a benefit over no ATFM; the benefits will be explained in Section 6.2.2.

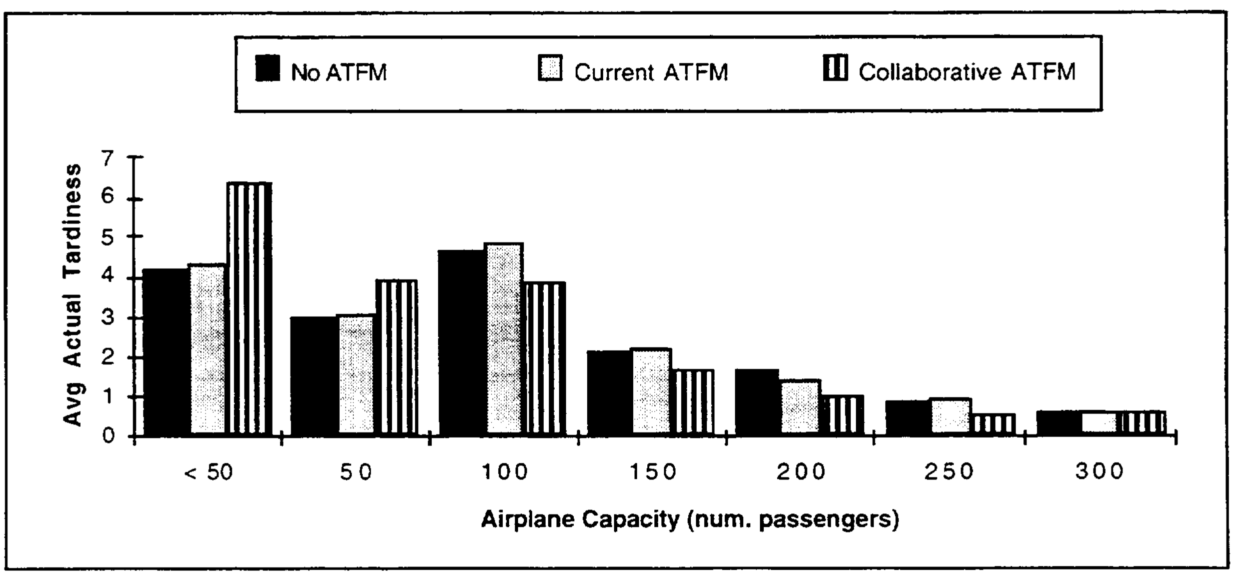

Figure 6.2.1-1: Effect of ATFM in blue sky conditions

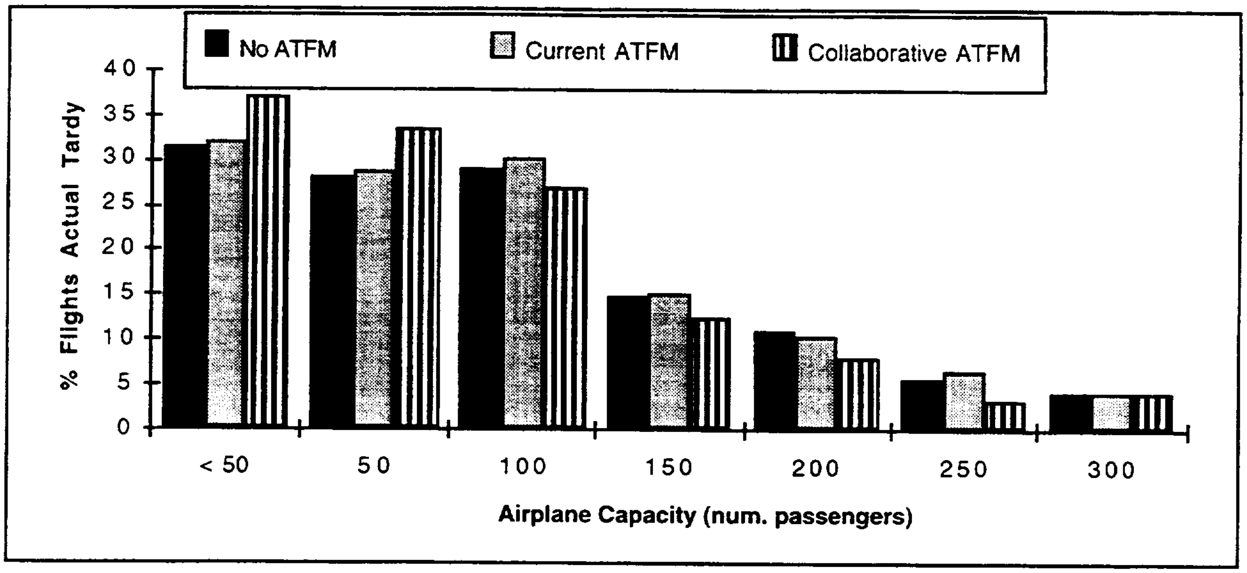

Figure 6.2.1-2: Benefit of ATFM in reducing % of tardy flights in blue sky conditions

6.2.2 Weather Front

In Figure 6.2.2-1, tardiness is broken into components: air delay, ground (gate) delay, and taxi delay. It illustrates a major benefit of ATFM: allowing unplanned delays to be taken as planned delays. In particular, air delay is reduced dramatically. This reduces aircraft density in the near-terminal area, making the system inherently safer and easier to control. In addition, airlines benefit by the reduced direct cost of operations. Results also indicate that taxi delay is reduced; this is a result of modeling taxi-in delays as a function of congestion.

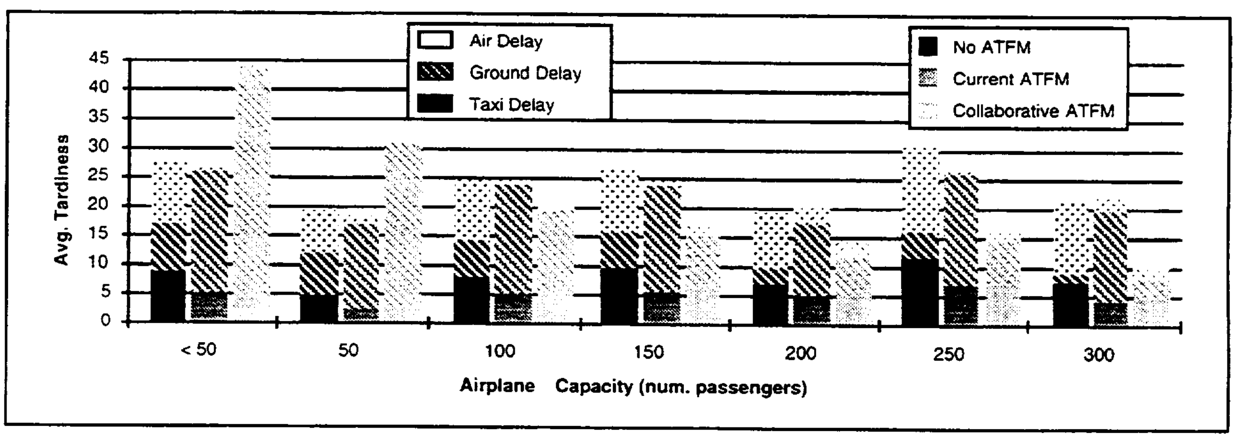

Figure 6.2.2-1: Benefit of ATFM in reducing average tardiness in a weather front

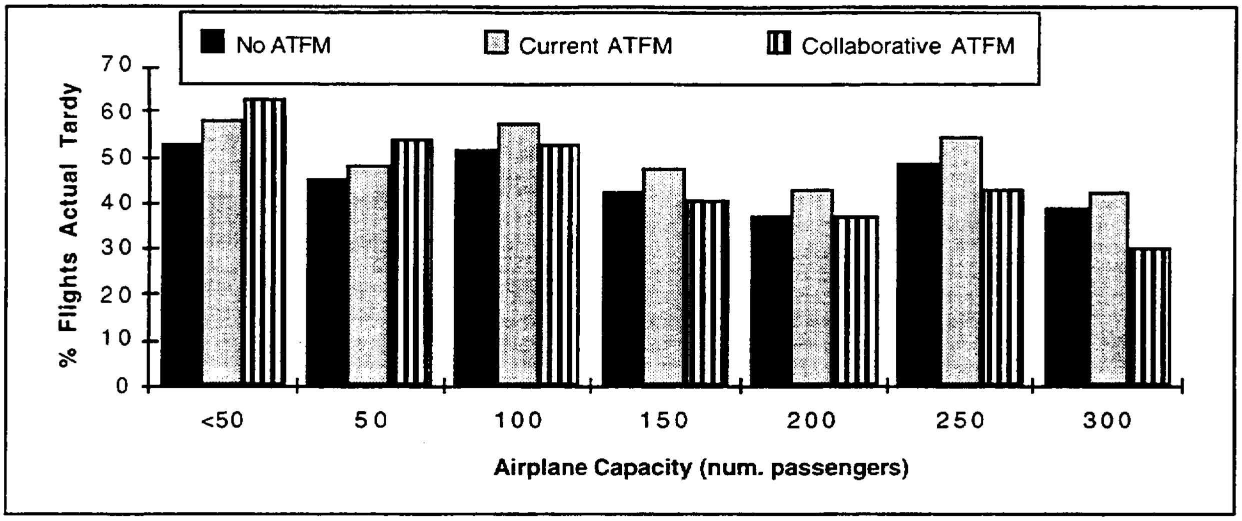

Figure 6.2.2-2: Benefit of ATFM in reducing % of tardy flights in a weather front

6.3 EFFECT OF COLLABORATIVE ATFM

This section discusses the benefits of implementing CSA. Here, it is assumed that all airlines are trying to reduce tardiness on their larger aircraft that hold more revenue-producing passengers. This is only a proxy for airlines' actual priorities in meeting schedule to achieve their business objectives. Of course, other priority schemes could be applied, with individual airlines employing schemes appropriate to their individual objectives. More research on individual airline operations is needed to improve this model of airlines operational policy. Collaborative ATFM is compared to the current ATFM both with and without the use of free flight and CTAS, and for both the blue sky and weather front scenarios. For each of these situations, three plots are presented. The aggregate statistics for all airlines is shown, followed by the metrics for each of two specific airlines, referred to in this report as Airline XX and Airline YY. Note that Airline XX has no aircraft with a capacity of 50 passengers and Airline YY has no aircraft with a capacity of 300 passengers.

6.3.1 Without Free Flight and CTAS

6.3.1.2 Blue Sky

Figures 6.3.1.2-1, 6.3.1.2-2, and 6.3.1.2-3, demonstrate that collaborative ATFM enables the airlines to reduce tardiness on higher priority flights. In particular, for the priority scheme chosen for these examples, collaborative ATFM decreases tardiness on flights with more than 50 passengers at the expense of the lower capacity flights. The magnitude of tardiness, in this blue sky condition, in some cases is so small (~ 1 minute) as to be in the background noise.

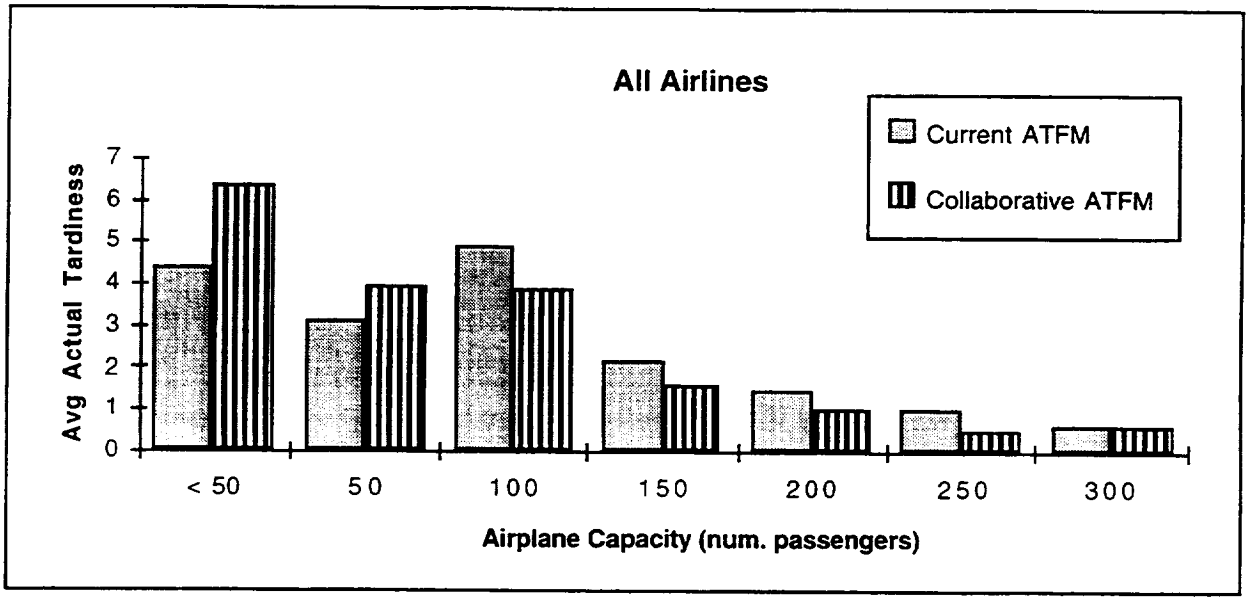

Figure 6.3.1.2-1: CSA reduces tardiness of higher priority flights in blue sky conditions

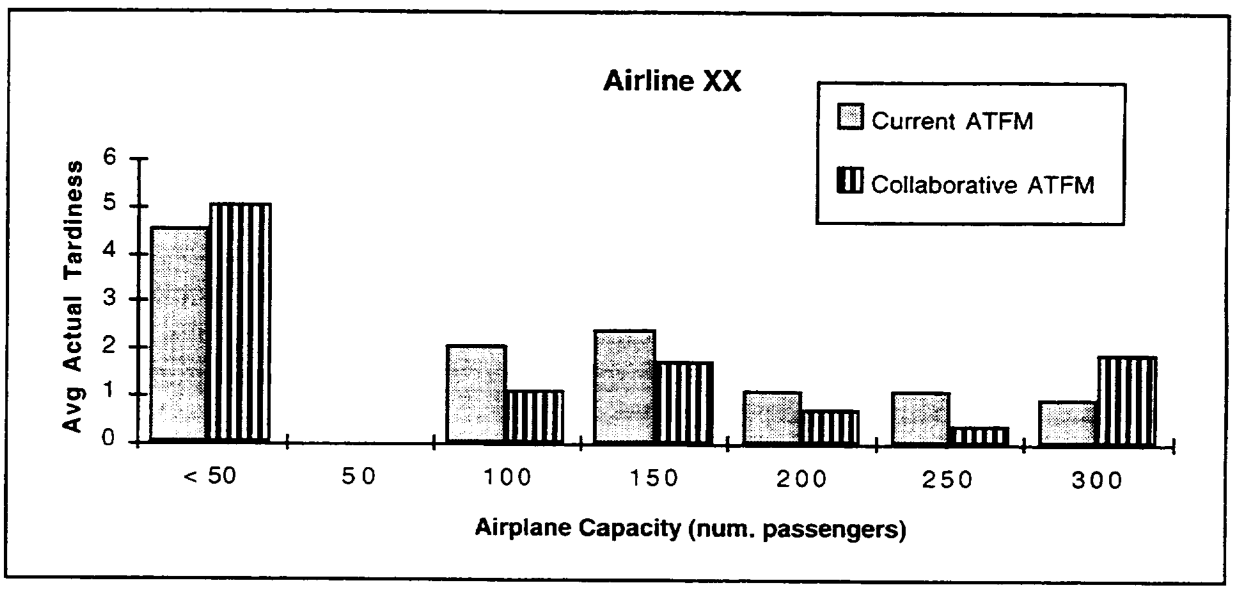

Figure 6.3.1.2-2: CSA reduces tardiness of Airline XX's high priority flights in blue sky

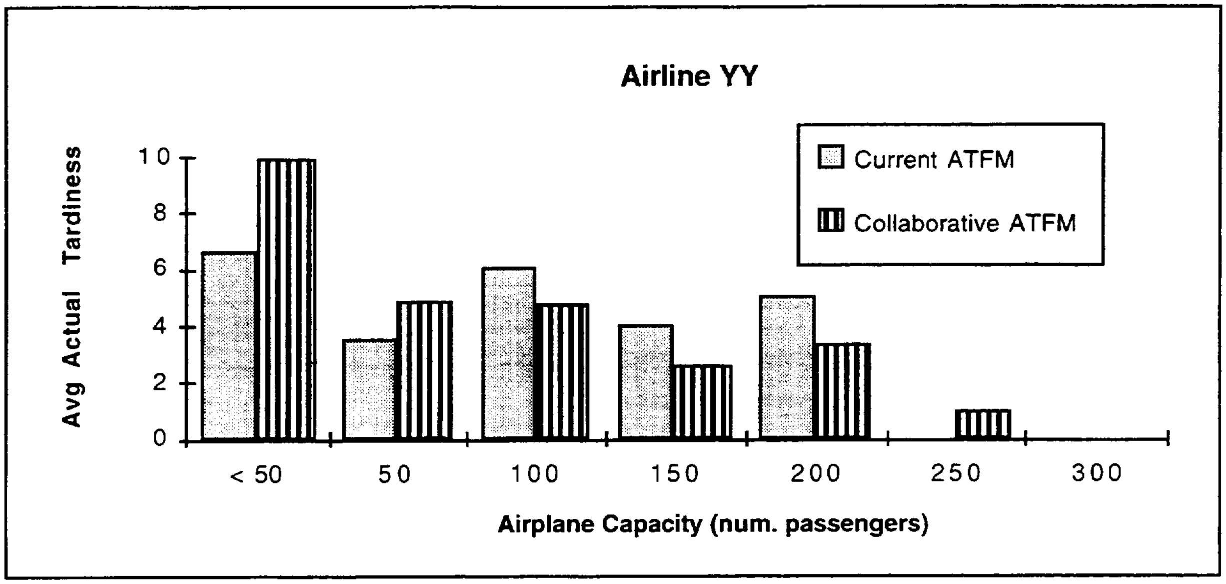

Figure 6.3.1.2-3: CSA reduces tardiness of Airline YY's high priority flights in blue sky

6.3.1.3 Weather Front

The effect of collaborative ATFM in shaping the delays to match its priority schedule is more dramatic here than under the blue sky scenario. The current ATFM algorithm results in delays spread fairly evenly across all classes of aircraft, as would be expected, since FSFS doesn't utilize airplane capacity information.

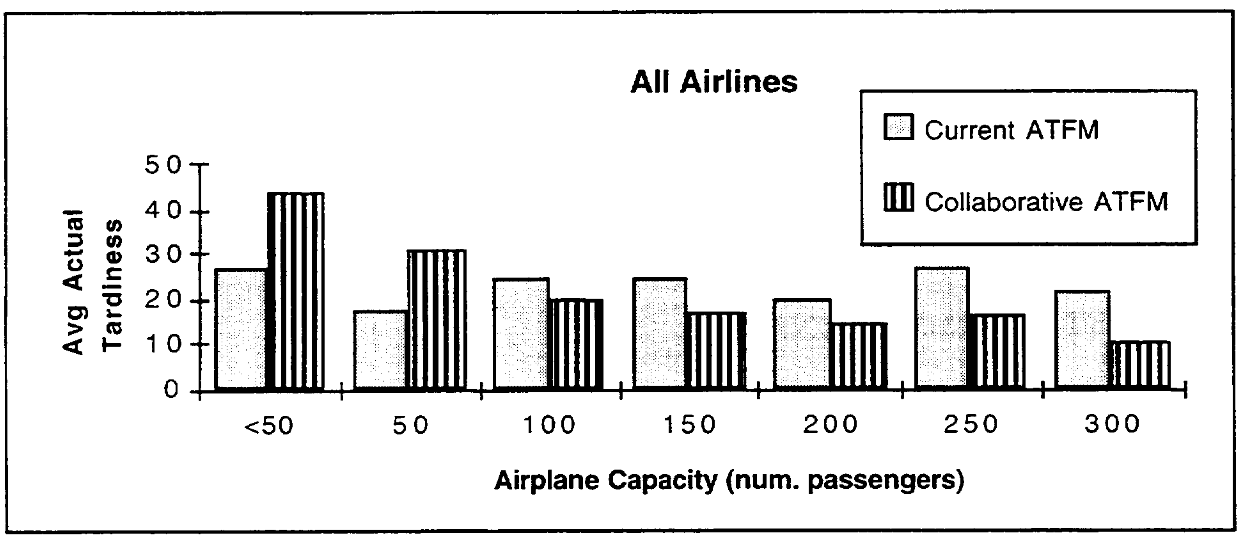

Figure 6.3.1.3-1: CSA reduces tardiness of higher priority flights in a weather front

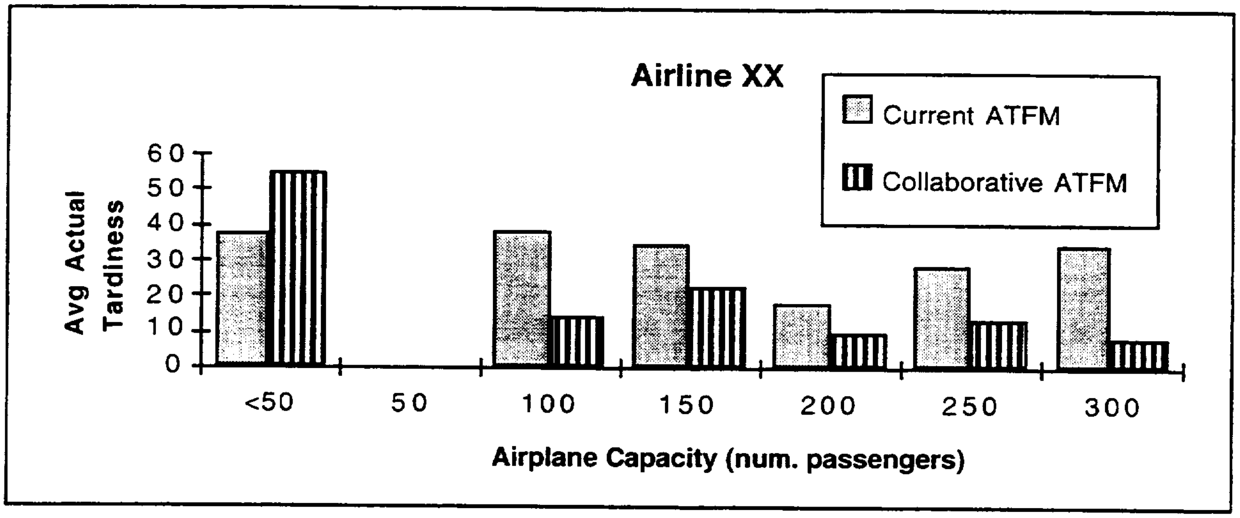

Figure 6.3.1.3-2: CSA reduces tardiness of Airline XX's high priority flights in a weather front

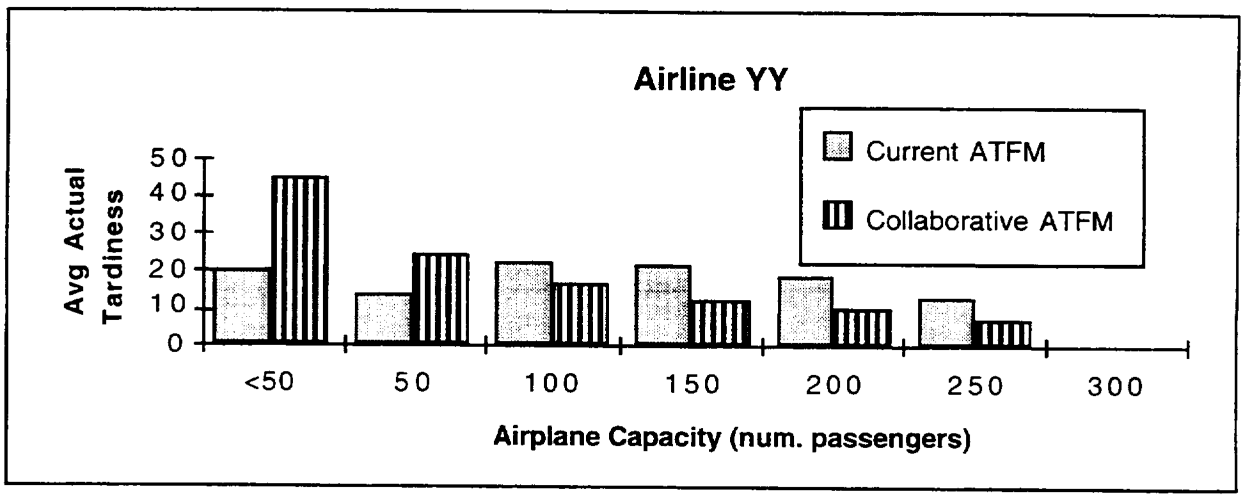

Figure 6.3.1.3-3: CSA reduces tardiness of Airline YY's high priority flights in a weather front

6.3.2 With Free Flight and CTAS

6.3.2.2 Blue Sky

With the added benefits from CTAS and UPR, the delays for the blue sky scenario in question are quite small; for example, Airline XX's largest average delay is 1.4 minutes. Nevertheless, collaborative ATFM shifts tardiness from flights with more than 50 passengers to the lower capacity flights.

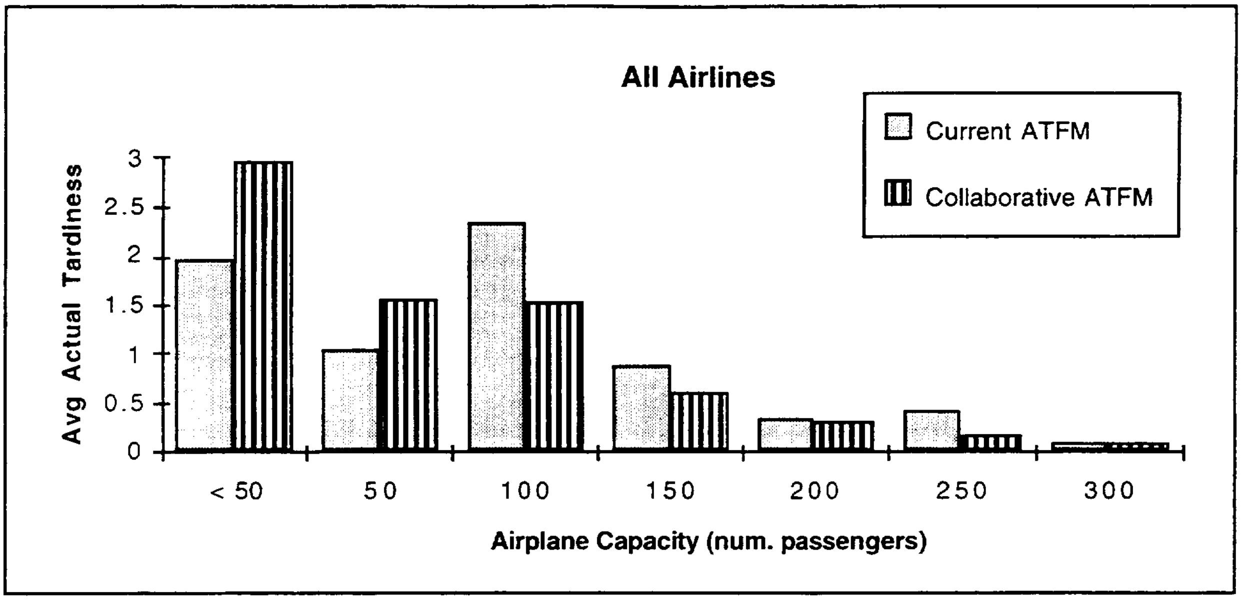

Figure 6.3.2.2-1: CSA reduces tardiness of higher priority flights in blue sky conditions CTAS and UPR are in effect

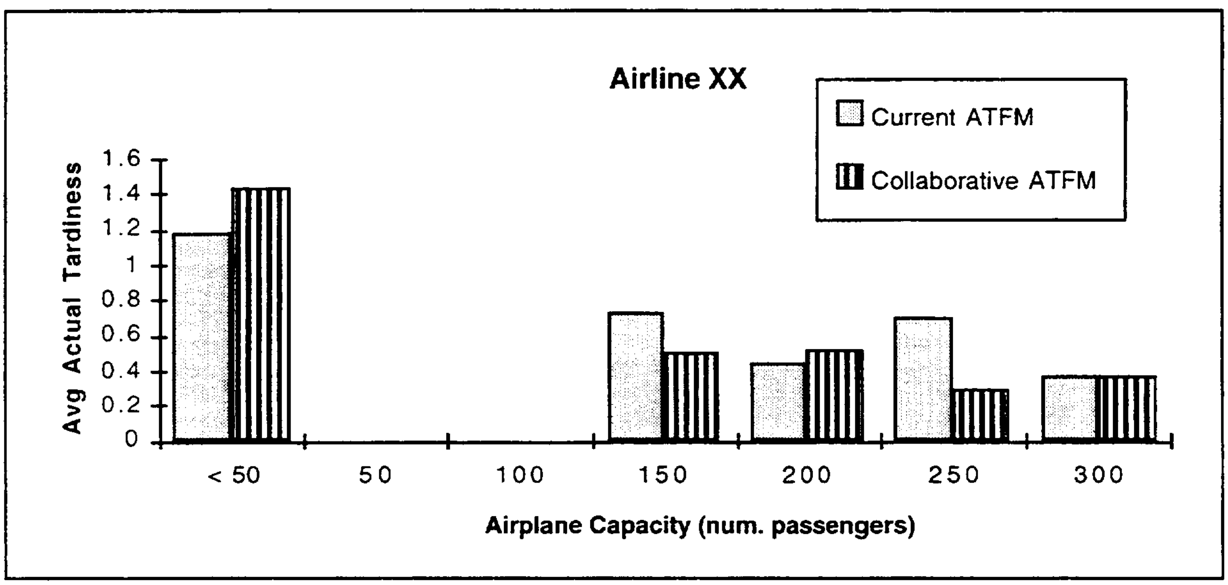

Figure 6.3.2.2-2: CSA reduces tardiness of Airline XX's high priority flights in blue sky conditions—CTAS and UPR are in effect.

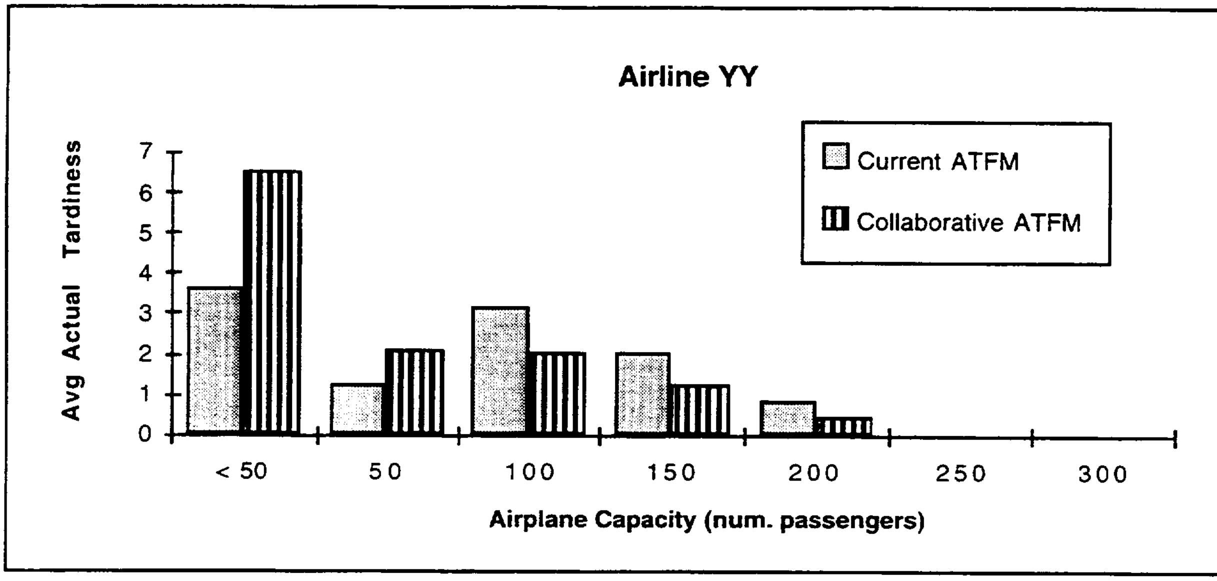

Figure 6.x CSA reduces tardiness of Airline YY's high priority flights in blue sky conditions—CTAS and UPR are in effect

6.3.2.3 Weather Front

Figures 6.3.2.3-1, 6.3.2.3-2, and 6.3.2.3-3, demonstrate that collaborative ATFM enables the airlines to reduce tardiness on higher priority flights. Collaborative ATFM decreases tardiness on flights with more than 50 passengers at the expense of the lower capacity flights.

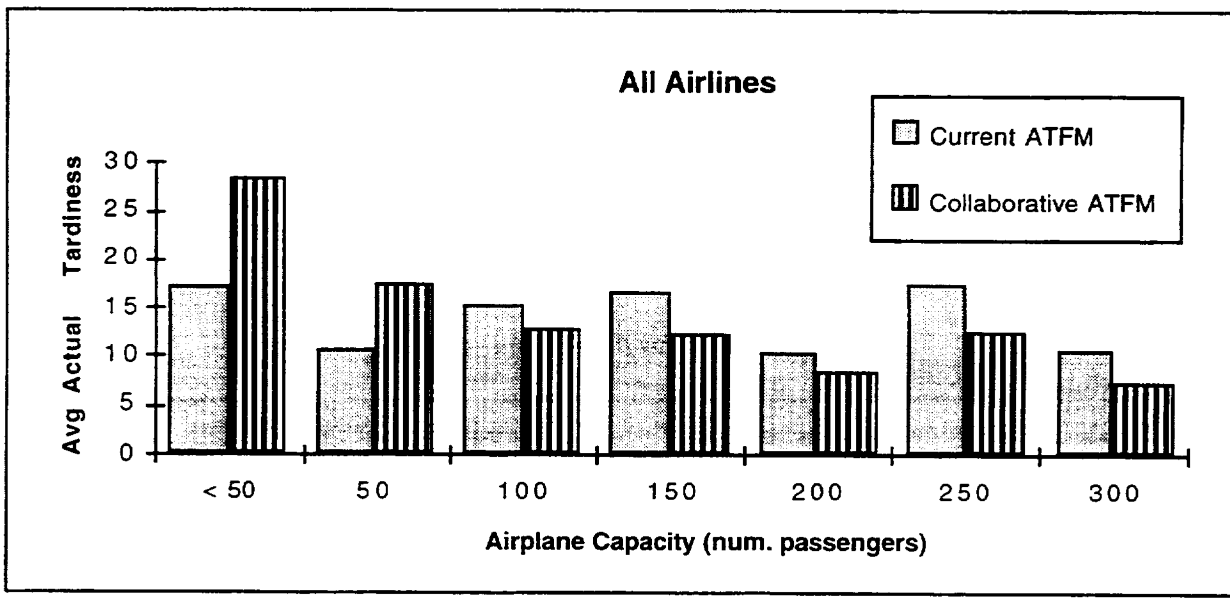

Figure 6.3.2.3-1: CSA reduces tardiness of higher priority flights with in a weather front CTAS and UPR are in effect

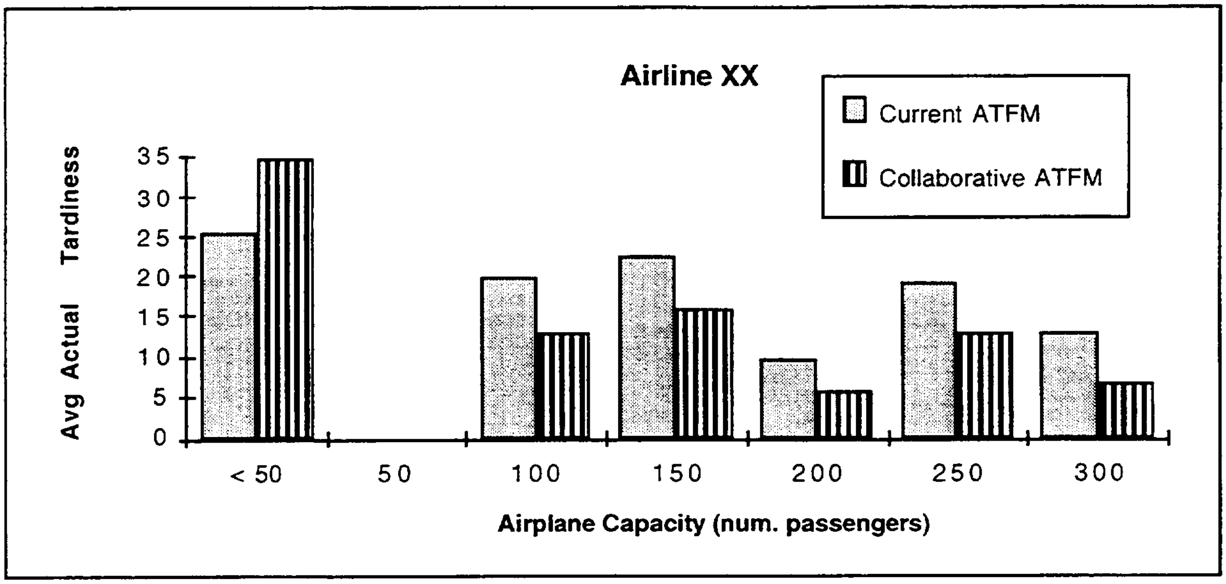

Figure 6.3.2.3-2: CSA reduces tardiness of Airline XX's high priority flights in a weather front—CTAS and UPR are in effect

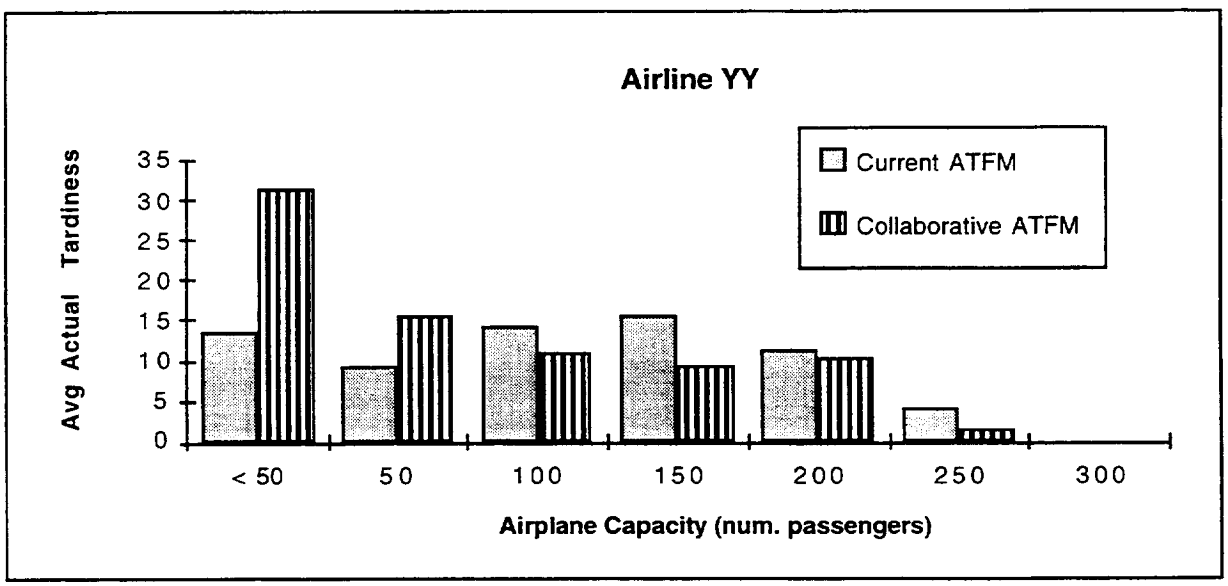

Figure 6.3.2.3-3: CSA reduces tardiness of Airline YY's high priority flights in a weather front—CTAS and UPR are in effect

6.4 Effect of Taxi Delays

This section examines the effect of taxi delays as a system component of tardiness. Figures 6.4-1 and 6.4-2 compare the inclusion of taxi delays as a component of the tardiness metric, illustrated in Figure 6.1. The *No ATFM* algorithm should provide the minimum for tardiness in the system–ATFM will redistribute delay (e.g., moving delay to the ground), but in general will not reduce it. The two figures demonstrate that our current models are probably attributing an unrealistic amount of taxi delays, causing the *No ATFM* tardiness totals to exceed those of *Current ATFM*. This result suggests the need for validation of our taxi delay models.

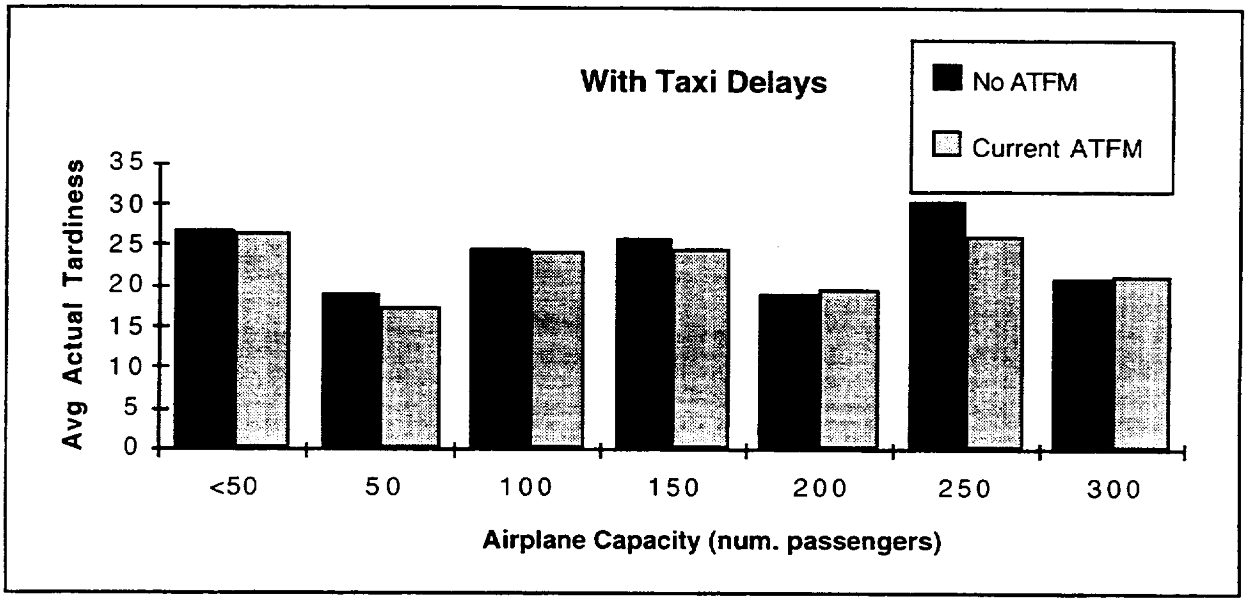

Figure 6.4-1: Taxi delays included in Tardiness measure during a weather front

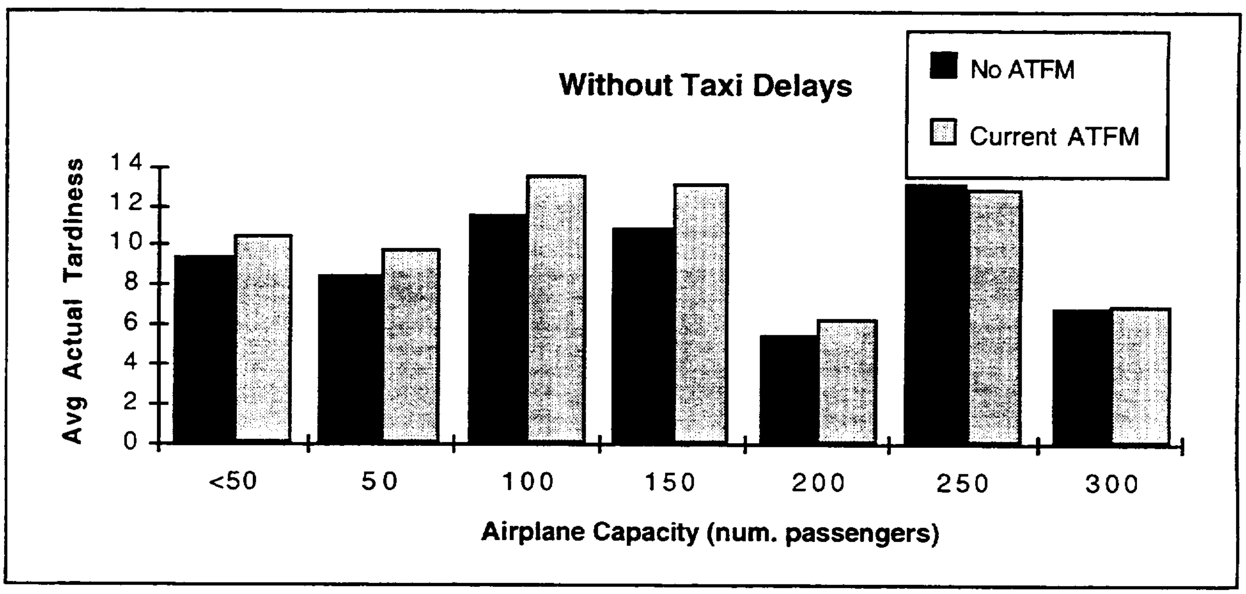

Figure 6.4-2: Taxi delays excluded from Tardiness measure during a weather front

6.5 Effect of Stochasticity on Transit Times

This section tested ASCENTs capability to add stochastic variance to transit times. A 10% variance on transit times was tested with no ATFM in order to identify the inherent delay in the system. Figure 6.5-1 shows the results of this run. Delays increased significantly for larger capacity aircraft–2 hours or more of tardiness was not uncommon on the higher capacity, long haul flights. This result suggests the need for validation of our transit time stochasticity model.

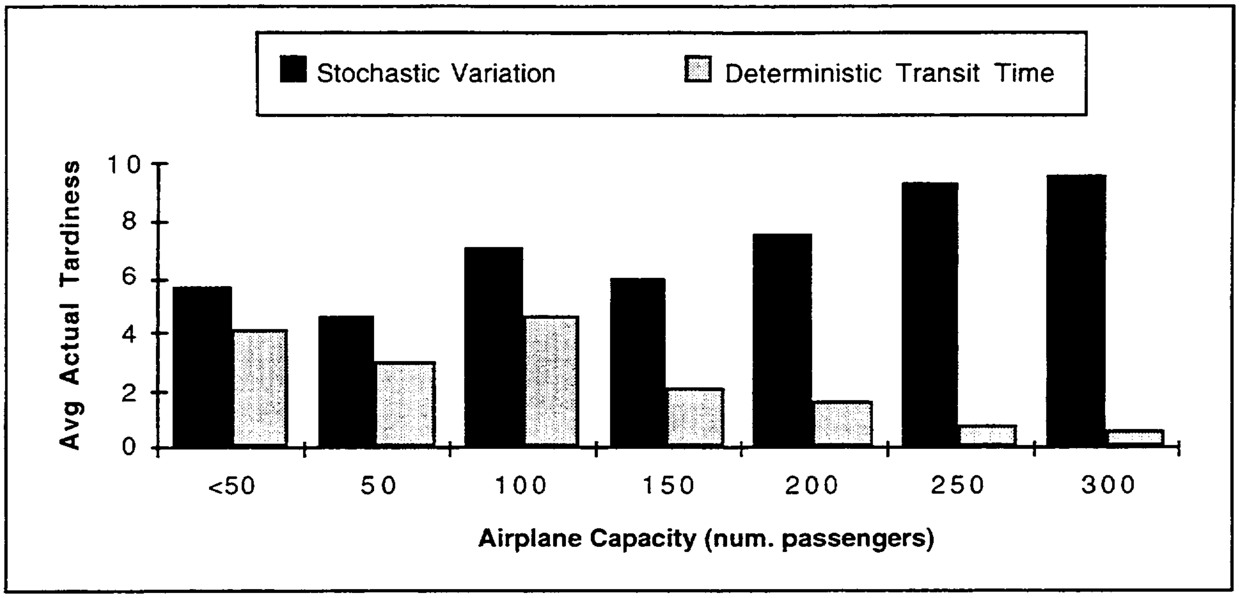

Figure 6.5-1: Effect of Stochastic Transit Times, No ATFM, Blue Sky

6.6 SUMMARY

1) We have selected a set of scenarios to illustrate the kind of results that can be generated for alternative ATFM strategies.

2) Benefits can be measured on a system-wide basis. A variety of metrics have been examined and since ASCENT provides detailed planning and realization information on each flight, other metrics could be examined, either directly in ASCENT or by post-processing the output.

3) Given other models of FAA and airline behavior, ASCENT could be enhanced to determine benefits of those as well.

4) **The Bottom Line:** these results are not meant to be conclusive or to suggest one approach over another—they are meant to be *representative* and we are ready to generate more when other candidates are proposed.

7. CONCLUSIONS

This report has presented a review of issues related to the evolution of the ATFM system in the United States toward more decentralized decision-making environments. A considerable range of conceivable alternative concepts has been broadly outlined in Section 4. For the short- to medium-term future, it would seem that one of these alternatives, described in more detail in Section 4.3 may be technically feasible, as well as consistent with expressed airline preferences and with the current emphasis on advancing free flight. It should be noted, however, that many open technical, procedural and operational issues need to be addressed with regard to such a partially decentralized system. A flexible simulation environment to support the evaluation and assessment of the benefits and costs stemming from such a concept is also needed; we have described here such an environment, along with a set of metrics that have been used to evaluate alternatives. We have illustrated the application of that testbed to the evaluation of a variety of free flight/collaborative decision-making alternatives to the current approach to ATFM. It is also clear that it is very difficult at this time to predict exactly how major airspace users, such as the airlines, would behave in decentralized ATFM environments. An approach to modeling and understanding some aspects of this behavior has been described here but this is a general area that will require extensive basic research over the next several years.

8. REFERENCES

Adams, Milton, S. Kolitz and A. Odoni (March 1996) "Addendum to Second Quarterly Technical Progress Report - Evolution Toward a Decentralized Air Traffic Flow Management System," Contract NAS2-14283(MJH), C.S. Draper Laboratory, Cambridge, Mass.

Andrews, John W. (1993) "Impact of Weather Event Uncertainty upon an Optimum Ground-Holding Strategy," *ATC Quarterly,* 1, 59-84.

Bertsimas, Dimitris and S. Stock (1994) "The Multi-Airport Flow Management Problem with En Route Capacities," Working Paper, Operations Research Center, MIT, Cambridge, Mass (to appear in *Operations Research).*

Brunetta, Lorenzo, G. Guastalla and L. Navazio (1995) "A New Approach for Solving the Multi-Airport Ground Holding Problem," Working Paper, University of Padua, Padua, Italy

DeArmon, James and A. R. Lacher (1996), "Aggregate Flow Directives as a Ground Delay Strategy: Concept Analysis Using Discrete Event Simulation," MITRE CAASD Working Paper, McLean, Virginia.

Erzberger, Heinz, T. Davis, and S. Green (1993) "Design of Center-TRACON Automation System," Proceedings AGARD Conference on Machine Intelligence in Air Traffic Management, Berlin, Germany.

Federal Aviation Administration (April 1996), FAA Free Flight Action Plan, Washington, DC.

Federal Aviation Administration (1975), *ATC Performance Measurement System at Major Airports*, FAA Technical Report, November 1975..

Gilbo, E. P. (1993), *Airport Capacity: Representation, Estimation, Optimization*, IEEE Transactions on Control Systems Technology, Vol.1, No. 3, September 1993.

Hocker, Guy A. (1994) *Airport Demand and Capacity Modeling for Flow Management Analysis,* Master's Thesis in Operations Research, MIT, Cambridge, Mass.

Milner, Joseph M. (1995) *Dynamic Slot Allocation with Airline Participation*, Ph.D. Dissertation, MIT, Cambridge, Mass.

Odoni, A. R. (1987) "The Flow Management Problem in Air Traffic Control", in *Flow Control of Congested Networks*, A. R. Odoni, Lucio Bianco and Giorgio Szego (eds.) Springer-Verlag, Berlin.

Richetta, O. and A. R. Odoni (1993) "Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control", *Transportation Science, 27*, pp. 228-238.

Robinson, J. D. (1992) *A Simulation Testbed for Flow Management Analysis in Air Traffic Control,* Master's Thesis in Operations Research, MIT, Cambridge, Mass.

RTCA (1994) *Report of the RTCA Board of Directors' Select Committee on Free Flight,* Washington, DC.

Swedish, W. J., (1981) Upgraded FAA Airfield Capacity Model, The MITRE Corporation Technical Report, MTR-81W16, 1981.

Vasquez-Marquez, Alberto. *American Airlines Arrival Slot Allocation System (ASAS).* Interfaces, 21 (1), 1991.

Wambsganss, Michael C. (1995) "Collaborative Traffic Flow Management," Working Paper, Metron Corporation, Washington, DC.

Wingrove, E.W., Lee, D. A., Kostiuk, P. F., and Hemm, R. V., (1995) *Estimating the Effects of the Terminal Area Productivity Program*, Logistic Management Institute, December 1995.

Yu, Jung (1996) *Airport Capacity and Regional Weather Modeling,* Master's Thesis in Operations Research, MIT, Cambridge, Mass.

9. APPENDIX

The following is a presentation of project research prepared for a site visit by NASA personnel during June, 1996.

Evolution Toward a Decentralized Traffic Flow Management System

June 21, 1996

NASA AATT Program

The Charles Stark Draper Laboratory, Inc.

and

The Massachusetts Institute of Technology

Schedule for day

9:00 AM	Schedule for day		
	Overview of project	Milt Adams, Stephan Kolitz	30 min
9:30	Introduction to CDM operational concept	Amedeo Odoni	15 min
9:45	Scenario generation	Brian Rexing, Stephan Kolitz	30 min
10:15	Tactical en route and transition area model	Bill Hall	30 min
10:45	break		15 min
11:00	Arrival slot scheduling heuristic	Amedeo Odoni, Mark Abramson	30 min
11:30	Airline tactical decision-making for slot utilization	Amedeo Odoni, Paul Carlson	30 min
12:00 N	Modeling the CDM operational concept	Stephan Kolitz	15 min
12:15	lunch		60 min
1:15	Testbed presentation and demo	Stephan Kolitz, Steve Nelson	45 min
2:00	Data needs	Stephan Kolitz	15 min
2:15	Re-routing in a capacitated network	Sarah Stock	30 min
2:45	break		15 min
3:00	Executive session		

Staffing

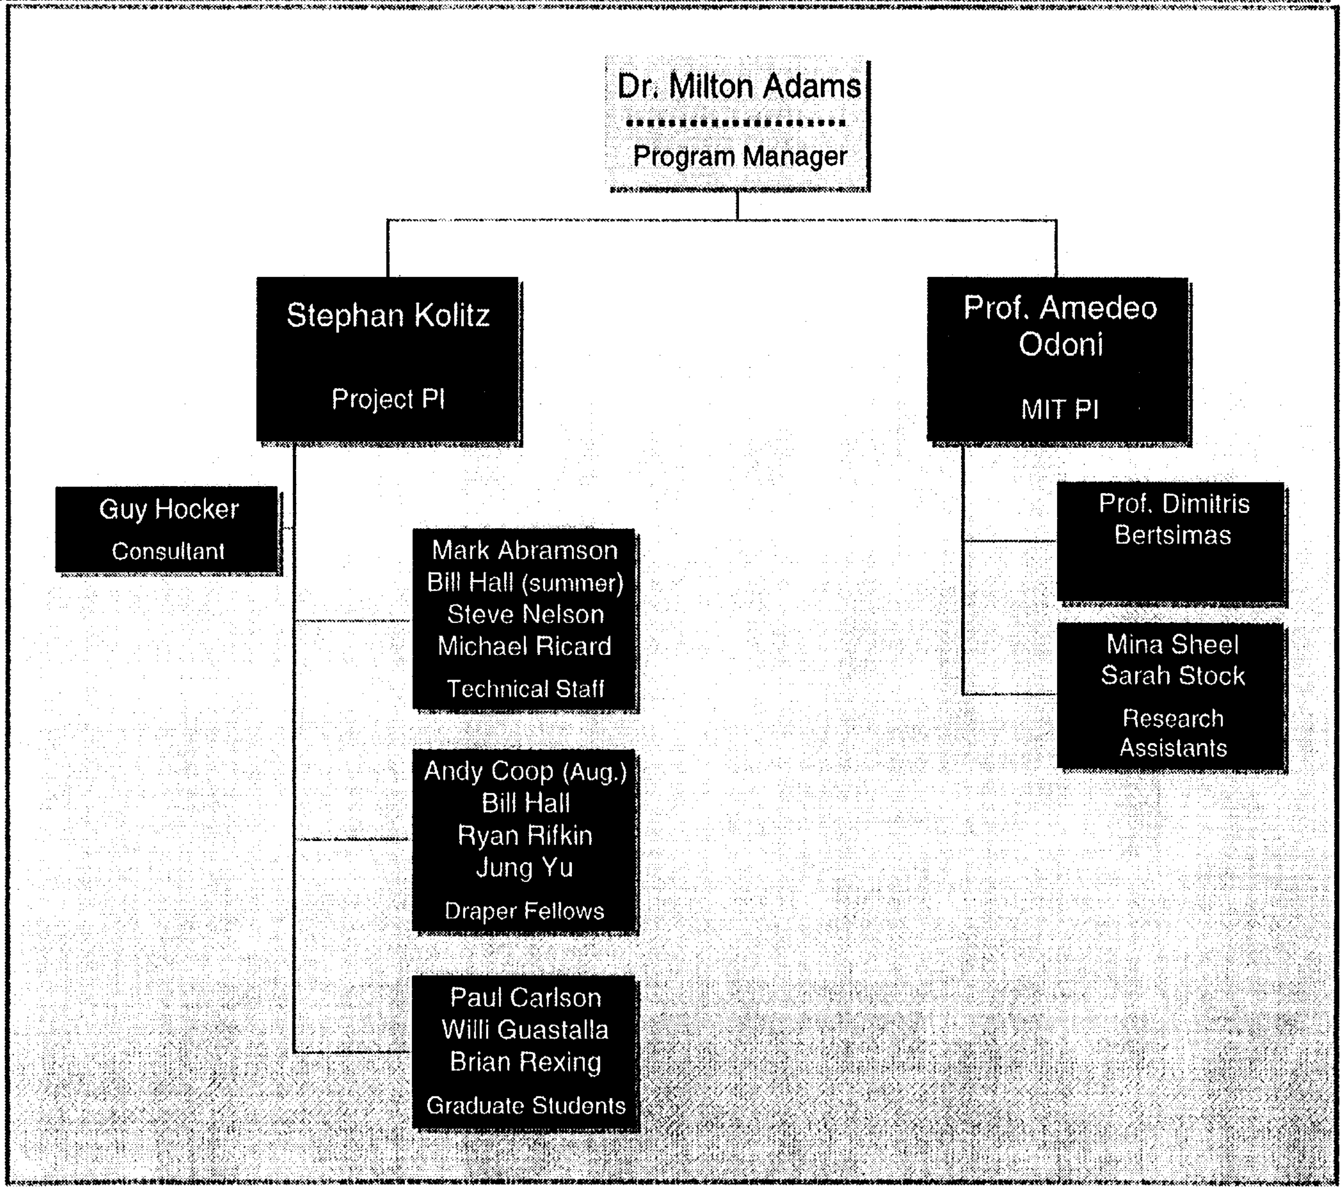
Dr. Milton Adams
Program Manager
Stephan Kolitz
Project PI
Prof. Amedeo Odoni
MIT PI
Guy Hocker
Consultant
Mark Abramson
Bill Hall (summer)
Steve Nelson
Michael Ricard
Technical Staff
Andy Coop (Aug.)
Bill Hall
Ryan Rifkin
Jung Yu
Draper Fellows
Paul Carlson
Willi Guastalla
Brian Rexing
Graduate Students
Prof. Dimitris Bertsimas
Mina Sheel
Sarah Stock
Research Assistants

Budget

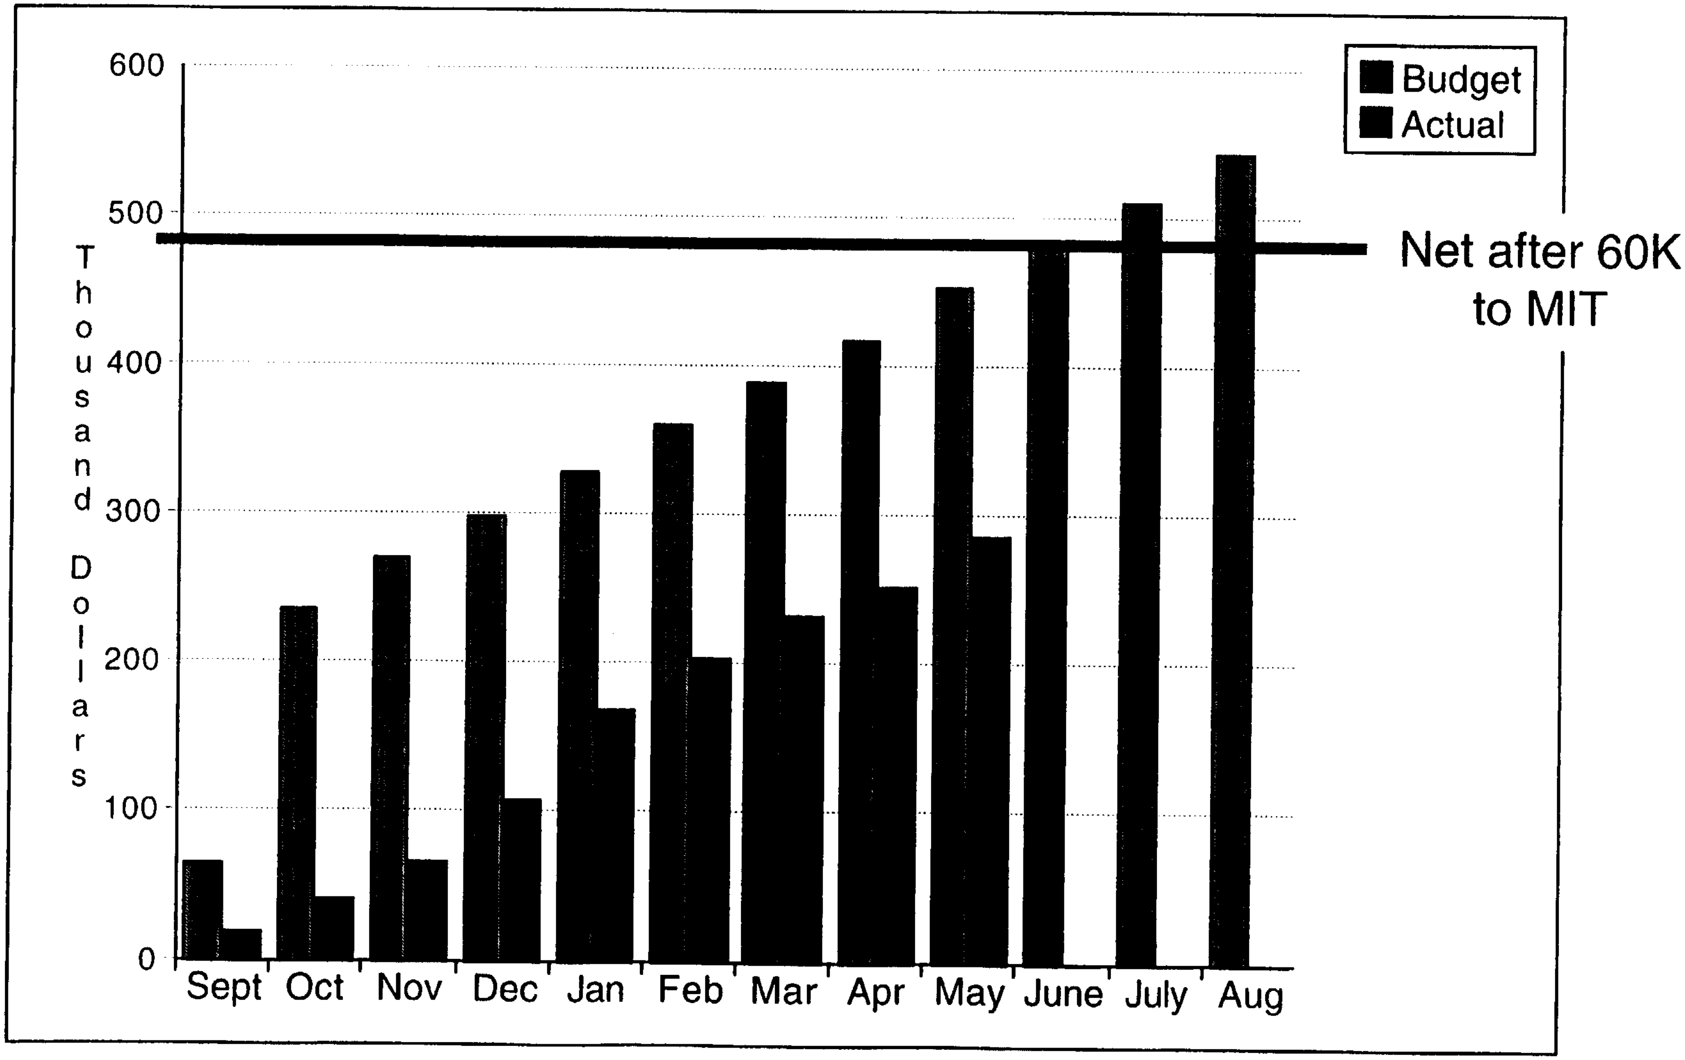

Budget
Actual
Net after 60K to MIT
600
500
400
300
200
100
0
Thousand Dollars
Sept
Oct
Nov
Dec
Jan
Feb
Mar
Apr
May
June
July
Aug

Overall Approach

➡ ***Objective*:**

Perform system-level analyses of collaborative tactical flow management under Free Flight

- **Describe a range of TFM concepts at varying degrees of decentralization**
- **Identify a viable medium-term concept**
- **Develop modeling capabilities for each element of the concept**
- **Integrate modeling capabilities into Draper ATFM Testbed**
- **Develop additional Testbed capabilities to facilitate evaluation**

Flow Management Functional Model

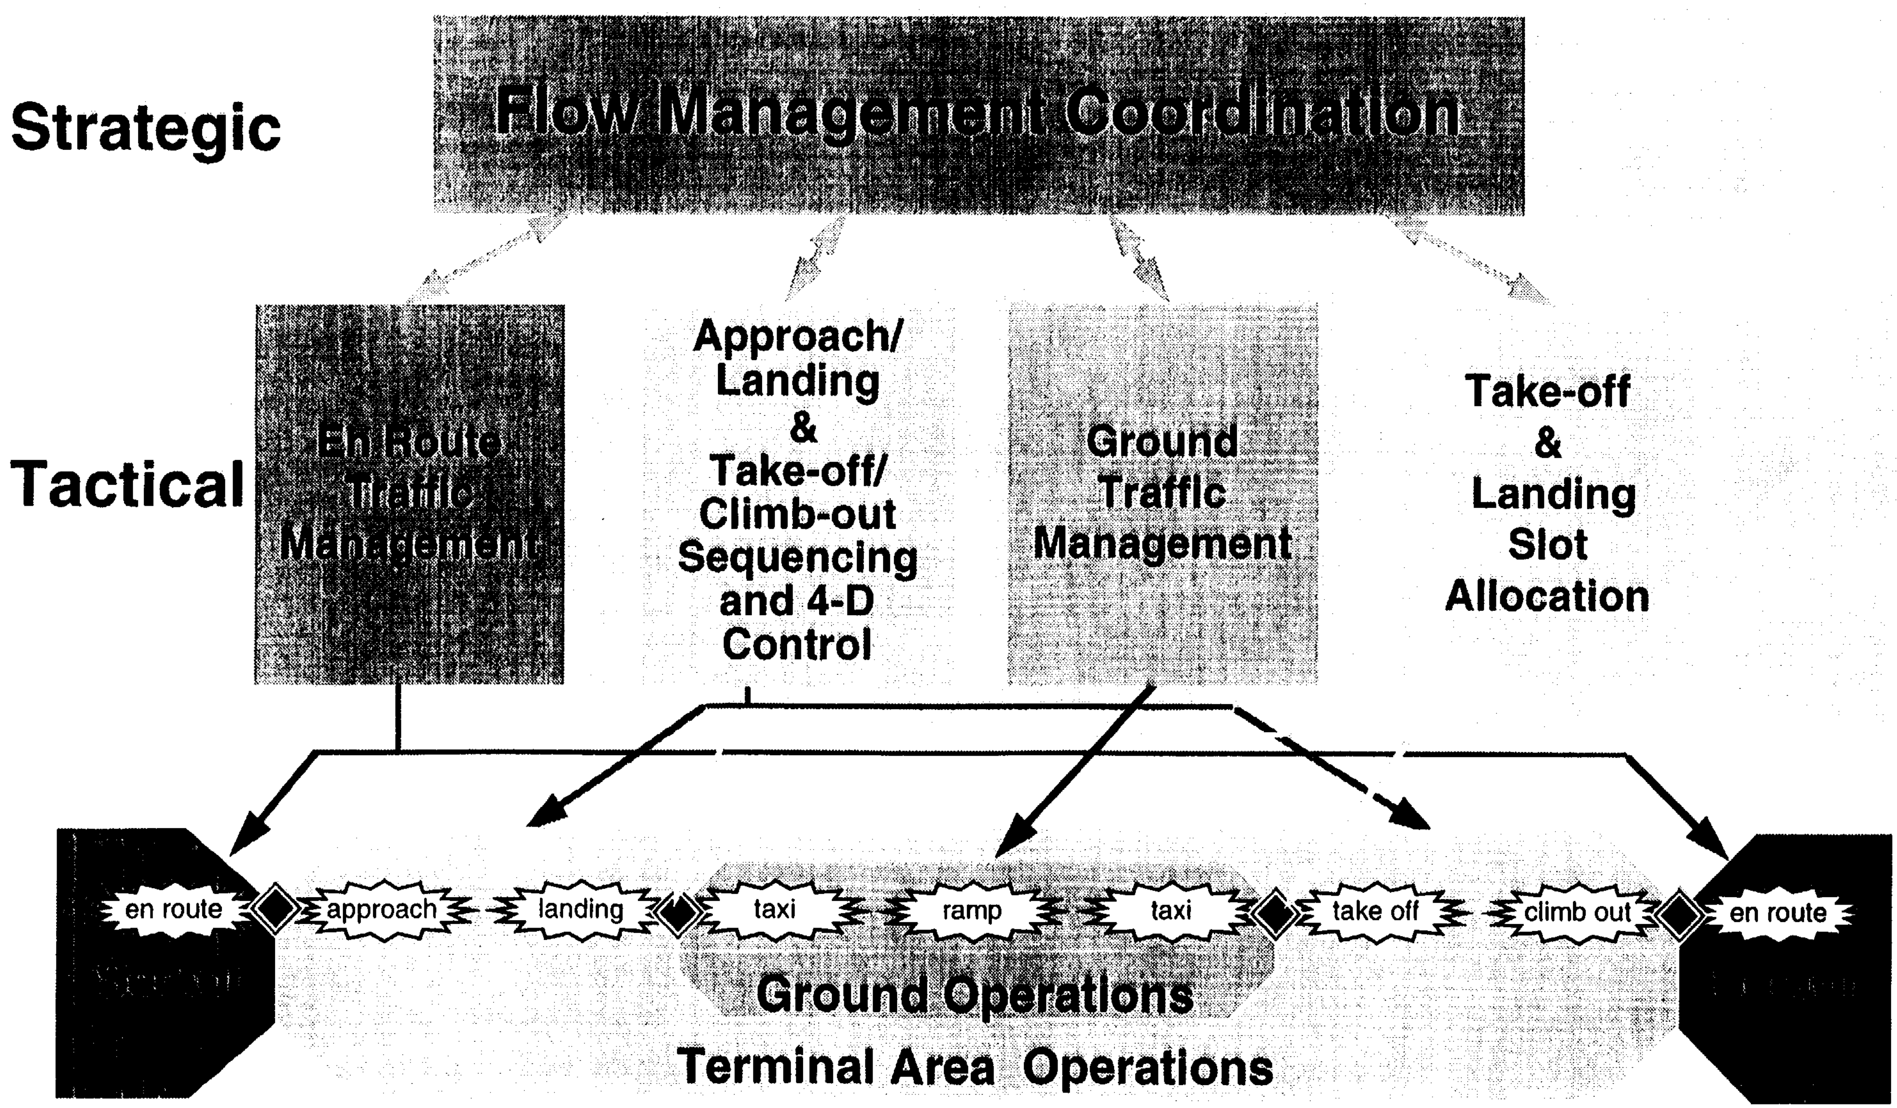

Network Flow

- Flow coordination across segments for an individual airport and its adjacent sectors is compounded when one considers that the ATFM system is a network
- Thus, system-wide (network-wide) measures must be addressed in evaluating collaborative tactical flow management
  - Focus on influence / coordination of flows at segment boundaries

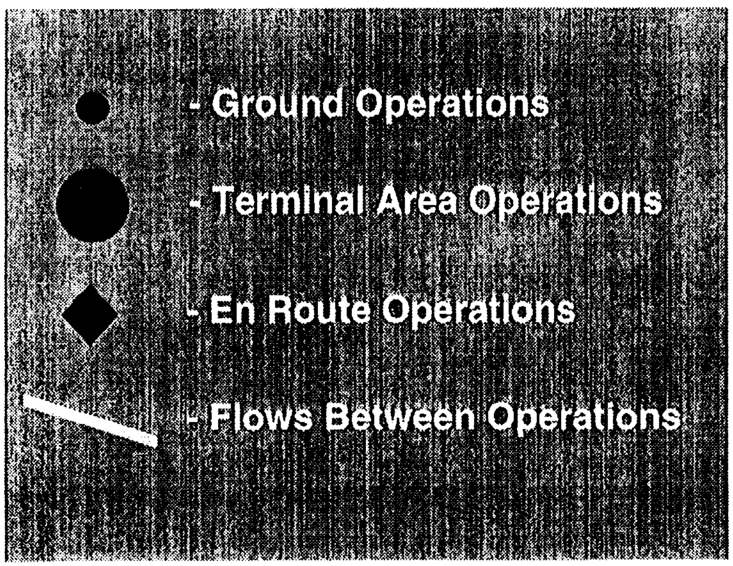

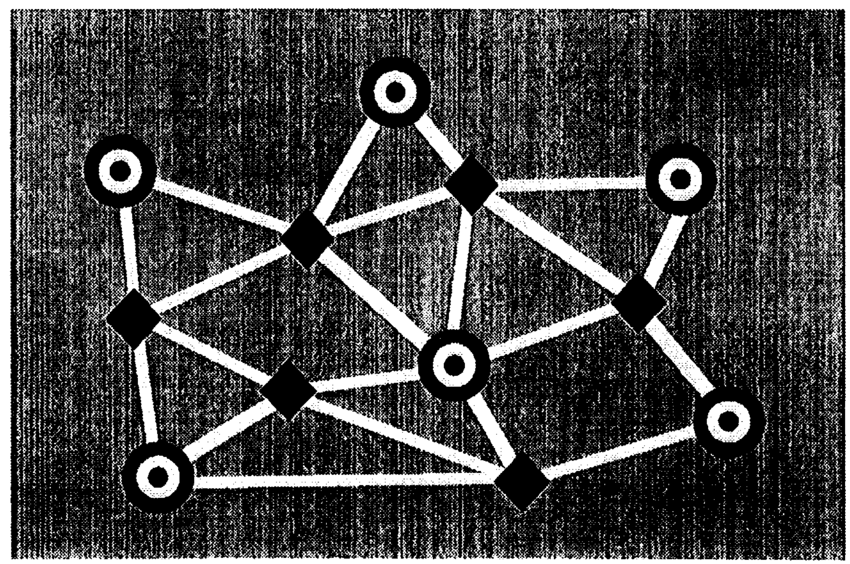

Scenario Generation

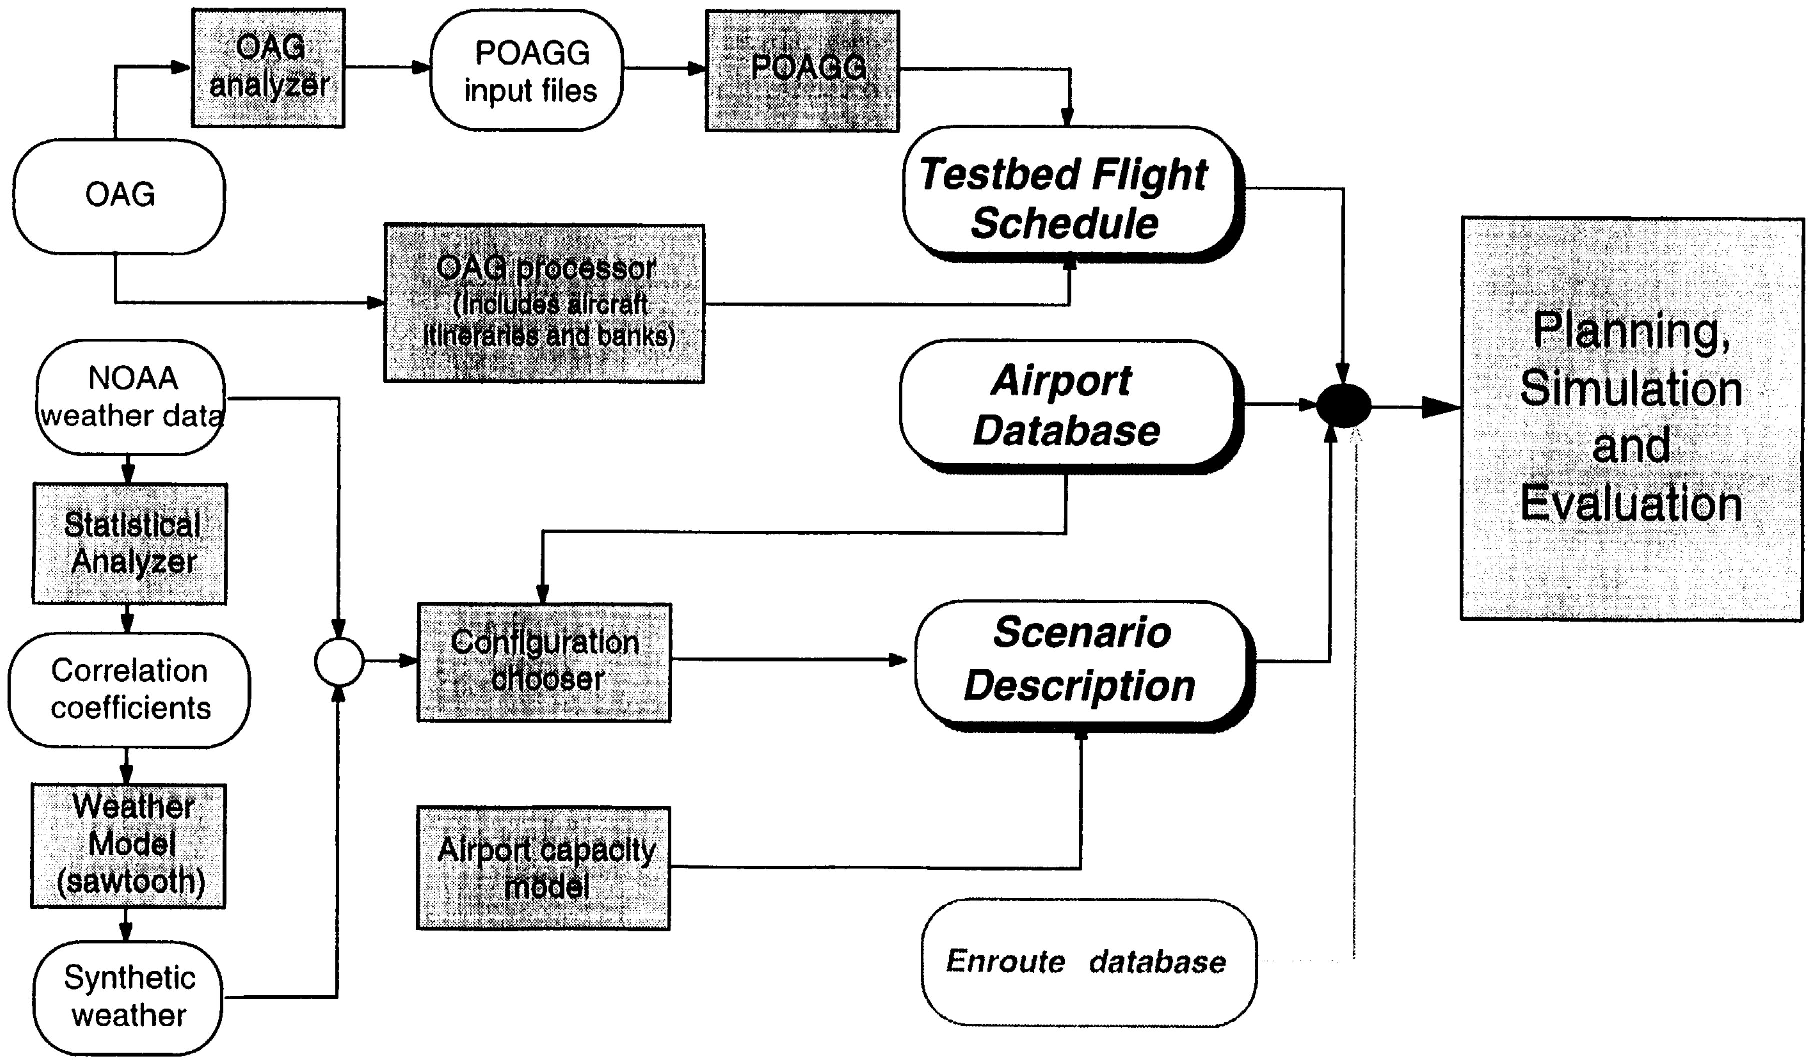

Planning, Simulation and Evaluation

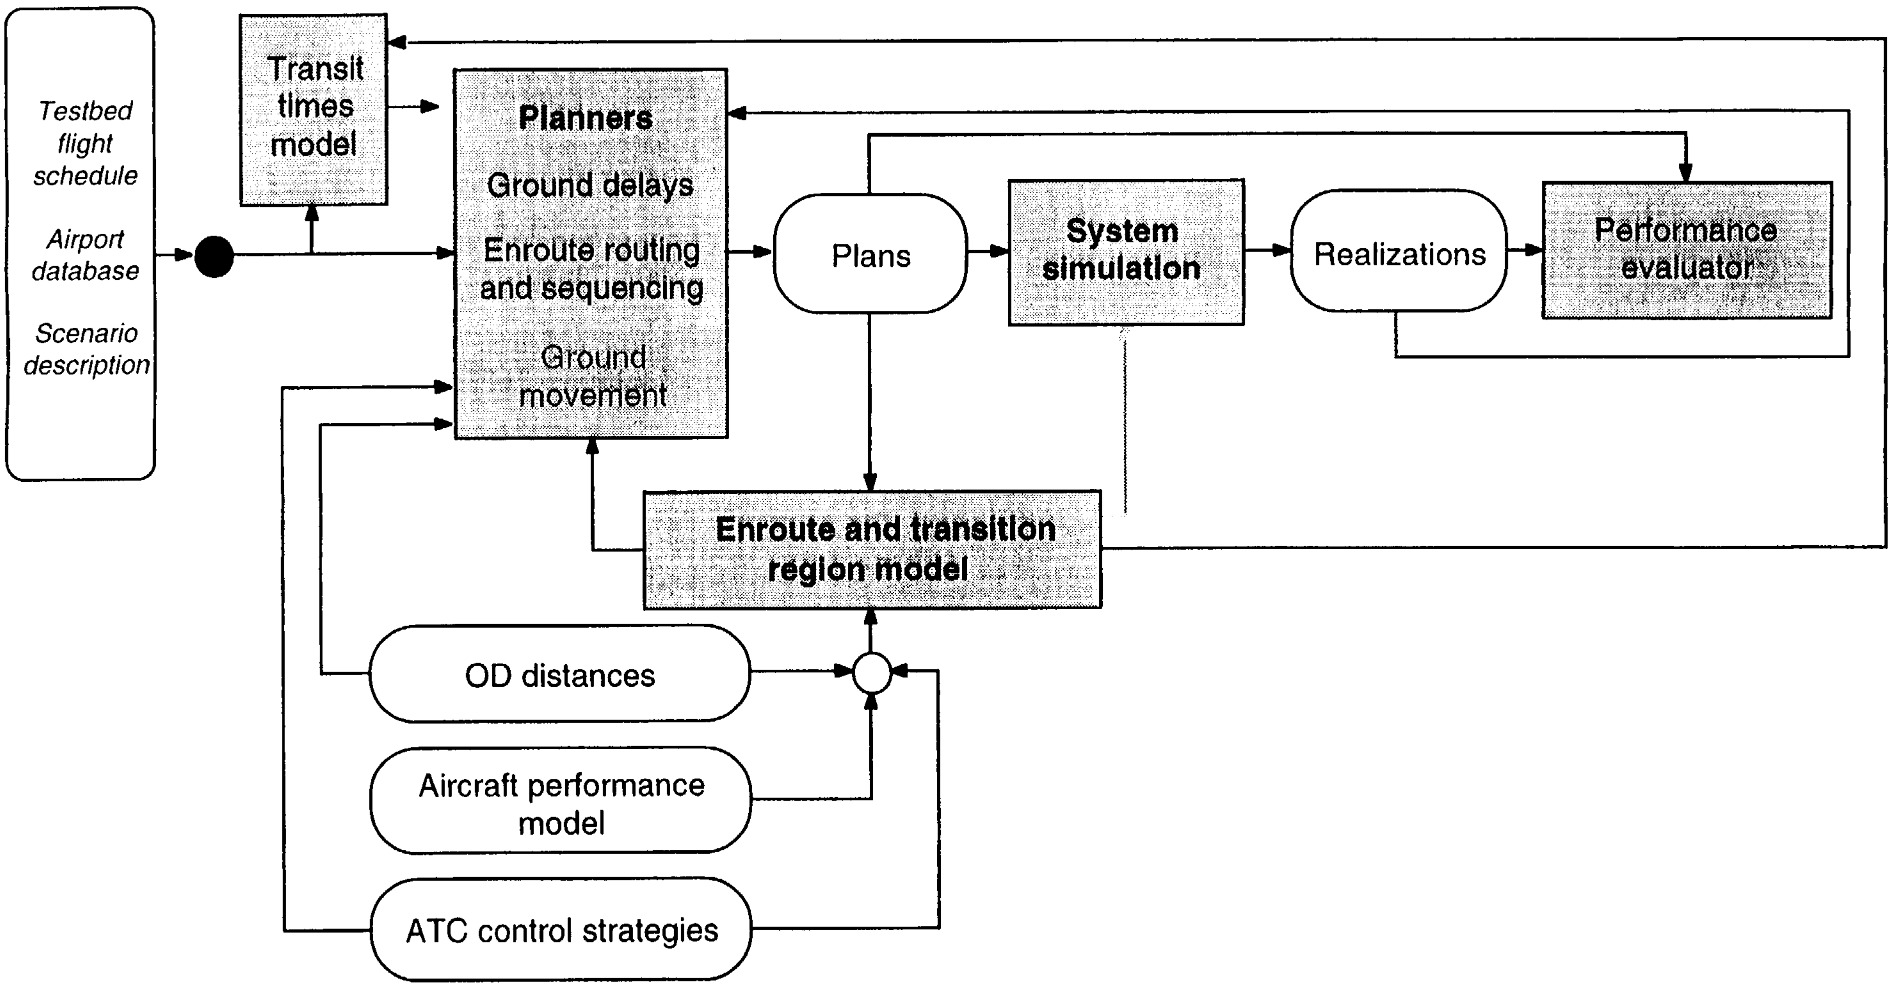

Descriptions of TFM Alternatives

Qualitative descriptions of TFM systems have been characterized by:

1. **Initial allocation of arrival slots among aircraft operators**
2. **Final assignment of arrival slots to individual flights**
3. **Assignment (if any) of departure slots to individual flights**
4. **En route flight planning and control**
5. **Transitional area, terminal area and airport surface flight planning and control**

Note: Dynamic, real-time decision-making is assumed throughout

Table of TFM Alternatives

	Allocation of Arrival Slots	Assignment of Arrival Slots to Individual Flights	Assignment of Departure Slots to Individual Flights	En Route Planning and Control	Transition Area, Terminal Area, Ground Movement Planning and Control
Centralized I	1a TFM system operator (FAA) allocates arrival slots to individual flights.	2a TFM system operator assigns slots; each airline may cancel and substitute flights	3a TFM system operator assigns departure slots to individual aircraft.	4a Airlines plan; TFM system operator controls.	
Partially Centralized (Current system)		2b Each airline suggests alternative assignment of the slots allocated to it (for its own flights only); TFM system operator approves or rejects.		4b Airlines plan; TFM system operator specifies a region in which airlines control their own aircraft; TFM system operator controls other regions and monitors globally for feasibility, conflicts.	
Partially Decentralized I	1b TFM system operator allocates sets of arrival slots to individual airlines.	2c Individual airlines allocate their own sets of slots among their own flights.	3b Airlines assign departure slots to individual aircraft; TFM system operator approves or rejects.		5a Airlines plan; TFM system operator controls.
Partially Decentralized II		2d Airlines may trade slots among themselves. Each individual airline allocates its own set of slots among its flights.			
Decentralized I	1c TFM system operator informs airlines of the legal safe capacities at potentially congested airports.	2e Airlines may bargain among themselves for the legal safe capacities. They may cancel or delay flights, follow the original schedule, etc. within their "purchased" slots.	3c Airlines assign departure slots to individual aircraft.		5b Airlines plan; they also specify each aircraft's heading directly after departure; TFM system operator can approves or rejects heading; TFM system operator controls everything else.
Decentralized II	1d TFM system operator informs airlines about anticipated availability of capacities at congested airports.	2e Airlines decide what they will do. They may cancel or delay flights, follow the original schedule, etc.		4c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.	5c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.

Current Area of Analysis

	Allocation of Arrival Slots	Assignment of Arrival Slots to Individual Flights	Assignment of Departure Slots to Individual Flights	En Route Planning and Control	Transition Area, Terminal Area, Ground Movement Planning and Control
Centralized I	1a TFM system operator (FAA) allocates arrival slots to individual flights.	2a TFM system operator assigns slots; each airline may cancel and substitute flights	3a TFM system operator assigns departure slots to individual aircraft.	4a Airlines plan; TFM system operator controls.	
Partially Centralized (Current system)		2b Each airline suggests alternative assignment of the slots allocated to it (for its own flights only); TFM system operator approves or rejects.			
Partially Decentralized I	1b TFM system operator allocates sets of arrival slots to individual airlines.	2c Individual airlines allocate their own sets of slots among their own flights.	3b Airlines assign departure slots to individual aircraft; TFM system operator approves or rejects.	4b Airlines plan; TFM system operator specifies a region in which airlines control their own aircraft; TFM system operator controls other regions and monitors globally for feasibility, conflicts.	5a Airlines plan; TFM system operator controls.
Partially Decentralized II		2d Airlines may trade slots among themselves. Each individual airline allocates its own set of slots among its flights.			
Decentralized I	1c TFM system operator informs airlines of the legal safe capacities at potentially congested airports.	2e Airlines may bargain among themselves for the legal safe capacities. They may cancel or delay flights, follow the original schedule, etc. within their "purchased" slots.	3c Airlines assign departure slots to individual aircraft.		5b Airlines plan; they also specify each aircraft's heading directly after departure; TFM system operator can approves or rejects heading; TFM system operator controls everything else.
Decentralized II	1d TFM system operator informs airlines about anticipated availability of capacities at congested airports.	2e Airlines decide what they will do. They may cancel or delay flights, follow the original schedule, etc.		4c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.	5c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.

A Viable Medium-Term Partially-Decentralized Scenario

RTCA Task Force 3; CDM TFM; MITRE; CTAS

1. 10-, 15- or 20-minute time intervals for allocation of arrival slots
2. When arrival capacity at airport A is scarce, available slots for arrivals at A are allocated to operators in FSFS way in each interval to ensure fairness (e.g., airline XYZ has 4 slots between 10:00 and 10:15)
3. Each operator is free to utilize its slots in each interval in the way deemed best
4. Each operator informs TFM of flight assignments to slots and of any slots left unused (due to flight cancellations)
5. TFM operator provides CTA's (controlled time of arrival) for each flight; where the point of "arrival" is the beginning of transitional area into terminal airspace of airport A; MAR is maintained

A Viable Medium-Term Partially-Decentralized Scenario II

6. **Little or no use of departure slot assignment (ECDT);**
   aircraft operator chooses preferred time of take-off to meet CTA

7. **"Free flight" (user-preferred routing) in large portions of en route airspace**

8. **Significant presence of automation aids in transitional/terminal area (e.g., CTAS) and on airport surface (e.g., SMA)**

9. **Monitoring by TFM for operator compliance with slot allocations; continuous monitoring for safety by ATC**

➨ **Numerous additional details to be specified; numerous possible variations on the above themes**

Airline Behavior

- **Most challenging aspects of evaluating decentralized TFM:**
  - Modeling airline behavior(s)
  - Quantifying benefits to airlines
- **Capture three types of airline considerations:**
  - Network effects (aircraft itineraries)
  - Schedule integrity (bank preservation)
  - Ground vs. airborne delay under uncertainty

Scenario Generation

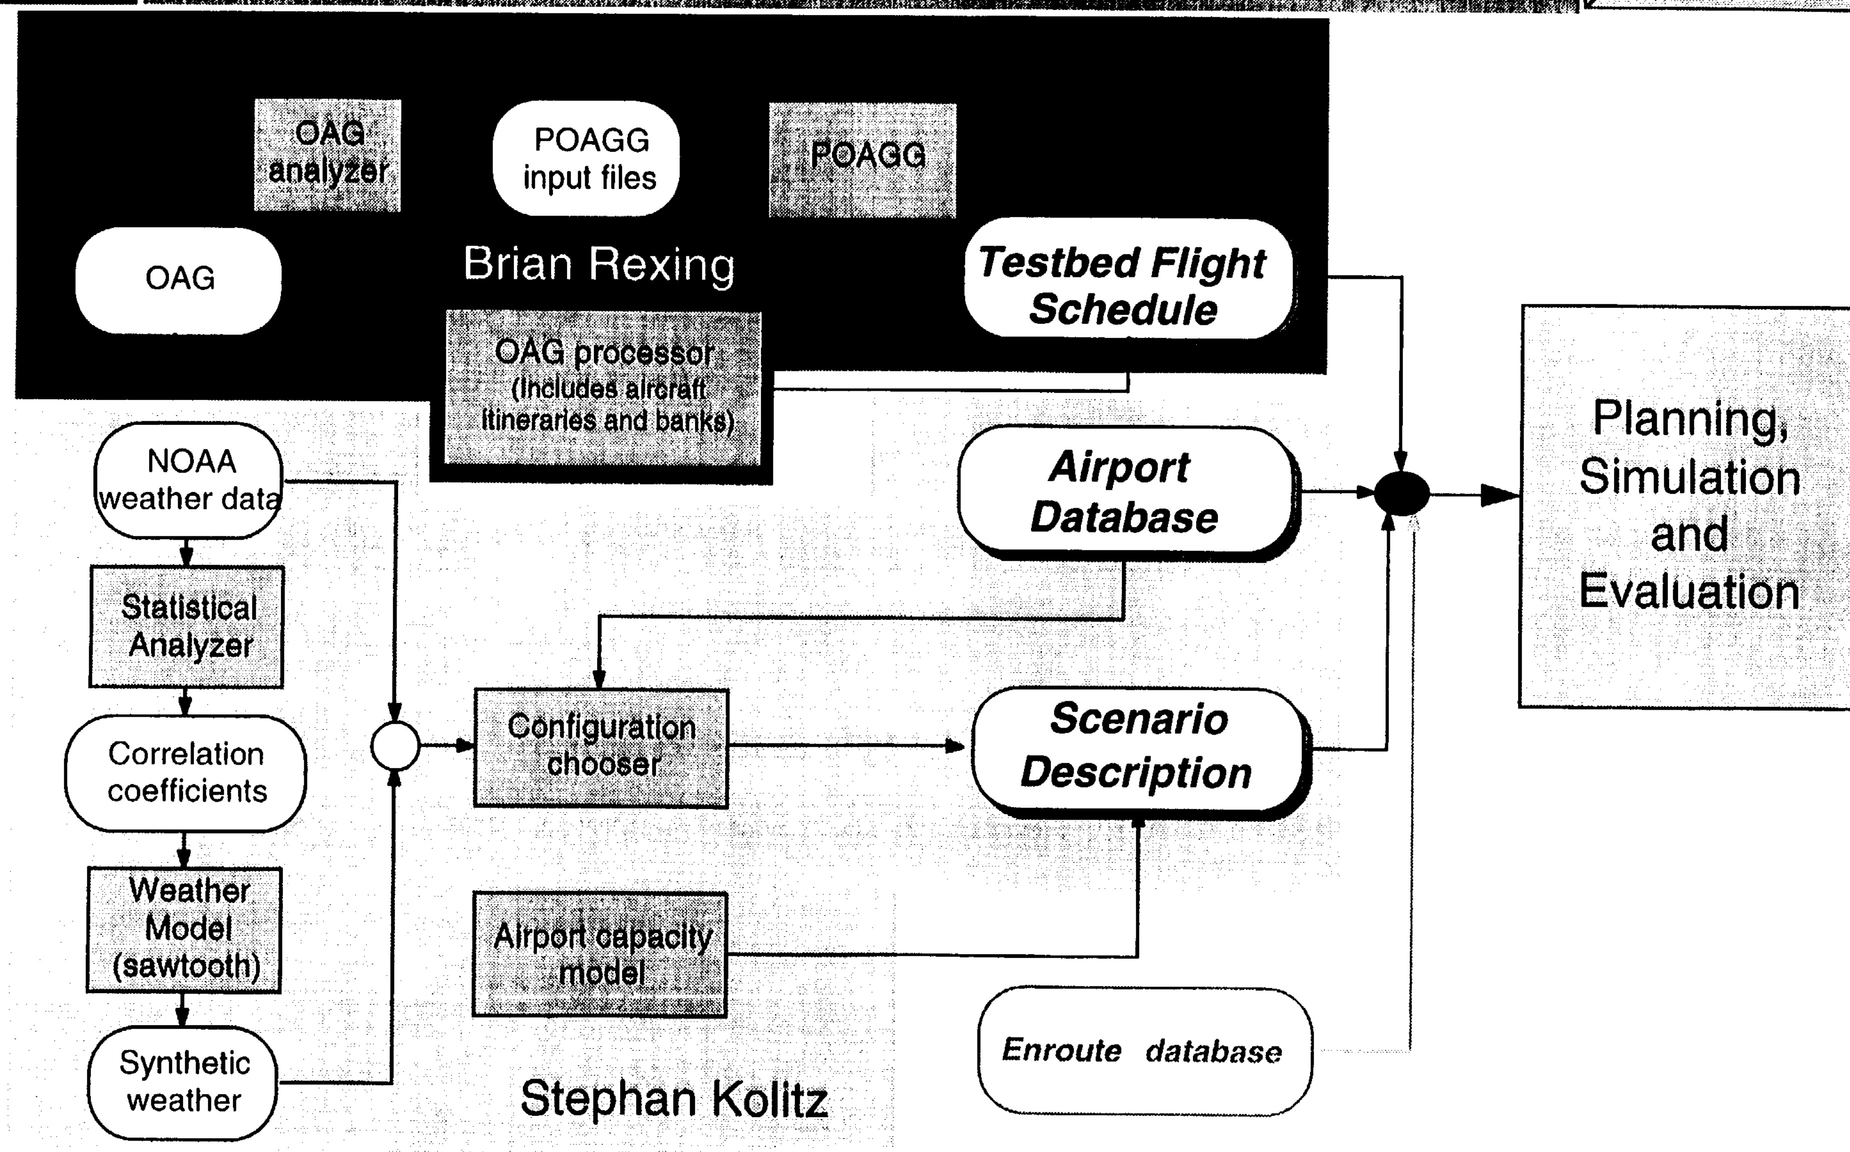

Purpose of POAGG Development

- Generate synthetic schedules that are statistically similar to OAG schedules
- Input for the Draper ATFM Testbed
- Model current as well as future scenarios
- Note: developed from a system, not airline, perspective

POAGG Input/Output

- **Input**
  - Number of capacitated airports
  - At each airport
    - Arrival distribution (relative frequencies of arrivals per period)
    - Carrier / aircraft %'s
  - Flight time matrix
  - Connection % matrix
  - Number of non-capacitated airports
  - Average flight time to each non-capacitated airport
  - Shuttle information
  - Bank information
- **Output consists of flight records that include**
  - Departure and arrival time
  - Airline, aircraft type
  - Previous and next flight of aircraft
  - Arrival and departure bank membership

POAGG Evaluation

- **Metrics**
  - Arrival distributions
  - Departure distributions
  - Connection percentages
- **Based on limited testing to date, POAGG generates accurate schedules from an overall system perspective**
  - Schedules indirectly reflect the airlines' perspective
- **A schedule directly developed from an airline's perspective:**
  - Market driven
  - Each airline has a network of airports
    - Hub (banks)
    - Non-hub
  - Based on routes, not flights
  - Limited number of aircraft, crews, maintenance facilities, etc.

Recent Testing

- Calculated POAGG input from OAG schedule
- Executed POAGG
- Compared statistics from POAGG schedules to OAG schedules
- Arrival distributions matched
  - Relative frequency of arrivals by 30 minute periods

Departures

Departures at the Five Capacitated Airports

Time	BOS		DCA		EWR		LGA		PIT		Total	
Period	OAG	POAGG	OAG	POAGG	OAG	POAGG	OAG	POAGG	OAG	POAGG	OAG	POAGG
00:00-	1	0	0	0	1	0	0	0	1	0	3	0
00:30-	0	0	0	2	0	0	0	0	0	0	0	2
01:00-	1	0	0	2	1	3	0	1	0	1	2	7
01:30-	0	0	0	0	0	1	0	0	0	1	0	2
02:00-	0	0	0	0	1	1	0	0	0	0	1	1
02:30-	0	0	0	0	0	0	0	1	0	0	0	1
03:00-	0	1	0	0	0	1	0	0	0	0	0	2
03:30-	0	0	0	0	0	0	0	0	0	0	0	0
04:00-	0	0	0	0	1	0	0	0	0	0	1	0
04:30-	0	1	0	0	0	2	0	0	0	1	0	4
05:00-	0	0	0	0	0	1	0	0	0	2	0	3
05:30-	0	4	0	3	2	6	0	0	1	2	3	15
06:00-	10	11	1	6	5	20	3	11	5	11	24	59
06:30-	14	12	10	14	17	14	25	10	2	15	68	65
07:00-	23	19	21	9	28	18	18	11	15	11	105	68
07:30-	20	24	9	11	19	13	17	10	7	7	72	65
08:00-	40	22	18	12	29	18	18	9	4	29	109	90
08:30-	22	20	8	14	33	18	20	22	34	22	117	96
09:00-	23	26	19	13	15	13	19	10	9	9	85	71
09:30-	16	13	6	9	10	8	12	14	19	27	63	71
10:00-	17	17	14	13	21	24	16	14	38	18	106	86

Departures (continued)

Departures at the Five Capacitated Airports

Time	BOS		DCA		EWR		LGA		PIT		Total	
Period	OAG	POAGG	OAG	POAGG	OAG	POAGG	OAG	POAGG	OAG	POAGG	OAG	POAGG
15:00-	21	27	9	15	14	8	17	16	22	23	83	89
15:30-	27	17	12	13	14	19	15	10	19	22	87	81
16:00-	29	30	11	10	30	27	16	13	7	25	93	105
16:30-	16	21	12	12	11	16	16	17	38	21	93	87
17:00-	31	23	18	13	19	21	14	19	11	14	93	90
17:30-	23	28	10	16	32	28	18	14	26	24	109	110
18:00-	27	29	13	10	27	20	20	21	32	24	119	104
18:30-	21	17	14	14	18	15	14	15	8	12	75	73
19:00-	25	24	17	15	7	24	13	17	2	1	64	81
19:30-	20	25	8	13	37	23	19	14	6	6	90	81
20:00-	20	26	18	12	20	22	13	17	28	34	99	111
20:30-	11	16	4	12	7	11	12	13	24	32	58	84
21:00-	9	17	8	14	2	15	8	17	1	20	28	83
21:30-	2	3	2	2	3	2	1	1	46	0	54	8
22:00-	12	2	1	2	2	1	2	1	13	1	30	7
22:30-	5	0	0	1	3	0	0	4	3	2	11	7
23:00-	1	0	0	0	1	2	0	2	1	0	3	4
23:30-	3	0	0	1	2	0	0	0	2	0	7	1
Total	654	641	373	387	577	569	481	462	565	553	2650	2612

Connections

Internal Flight Count

OAG

	BOS	DCA	EWR	LGA	PIT	XXX	Depart.
BOS	0	32	35	34	9	544	654
DCA	31	0	18	31	8	285	373
EWR	35	17	0	0	14	511	577
LGA	34	31	0	0	9	407	481
PIT	9	8	14	8	0	526	565
XXX	542	286	515	407	525	0	2275
Arriv.	651	374	582	480	565	2273	4925

POAGG

	BOS	DCA	EWR	LGA	PIT	XXX	Depart.
BOS	0	30	29	29	9	544	641
DCA	30	0	17	30	8	302	387
EWR	40	10	0	0	12	507	569
LGA	34	30	0	0	8	390	462
PIT	10	6	15	7	0	515	553
XXX	540	302	522	414	531	0	2309
Arriv.	654	378	583	480	568	2258	4921

Motivation for Weather Modeling

- **Increased demand for air resources**
  - 400 million passengers in 1990
  - 800 million passengers in 2000
- **Capacity problems cause significant delays in system**
  - 50 million hours of passenger delay in 1990 at ten busiest airports
- **Weather - most dominant factor in capacity**
  - FAA: 65% of all delay weather related
  - Low visibility: increased spacing intervals
  - Wind: eliminate configuration choices
- **Generation of synthetic weather useful in modeling of scenarios**
  - Synthetic weather: observations statistically identical to historical data

Weather Components

- **Cloud Ceiling**
  - Base altitude of lowest opaque cloud in the terminal area
- **Horizontal Visibility**
  - Visibility at an airport at ground level
- **Visibility and ceiling determine the flight rule conditions at an airport**
- **Wind determines what runways could be used**

Sawtooth Wave Model

- **Derived from validated USAF weather model**
- **Models spatial correlation**
  - E.g., Boston, New York, Washington DC
- **Models temporal correlation**
  - Weather at 10 AM is correlated with weather at 6 AM
- **Model cross-variable correlation**
  - Cloud ceiling and horizontal visibility are correlated
- **Limited to a region of the country, but can integrate multiple SWMs**

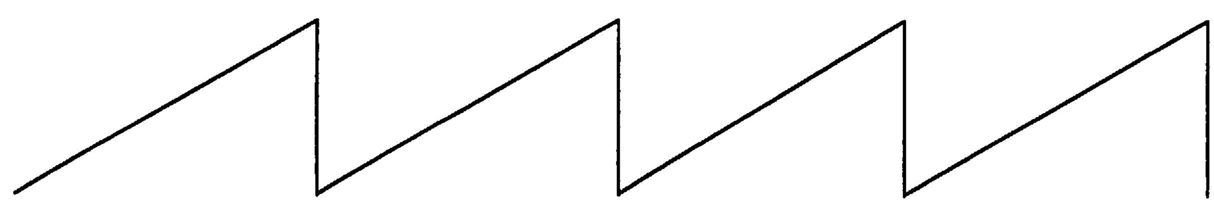

Capacity Models and Issues

- **Empirical Data Capacity Frontiers (EDCF)**
  - Data intensive
  - Low capacity estimates even though observations taken during peak periods

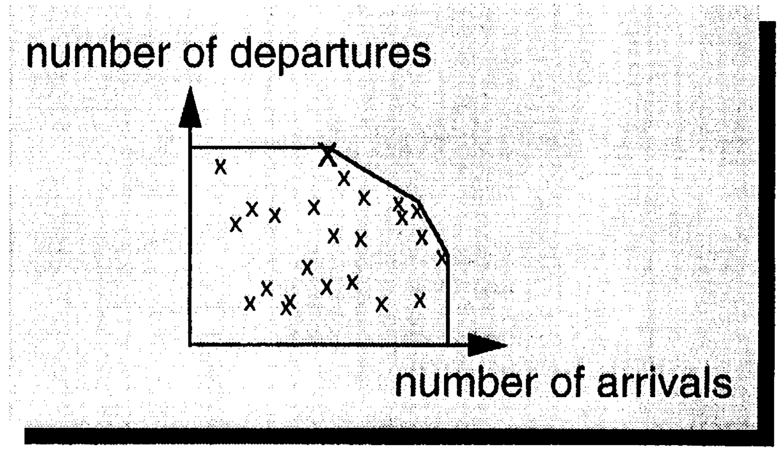

- **Engineered Performance Standards (EPS)**
  - Available
  - Not enough values on the frontier
- **FAA Airfield Capacity Model**
  - Difficult to generate accurate input values
    - E.g., means and variances of runway occupancy times and interarrival times
  - Several assumptions simplify model too much
    - Flight rule conditions do not affect capacity for certain configurations
    - Ignores capacity of an runway intersecting a pair of parallel runways
    - Airport-specific operations are ignored
- **LMI Airfield Capacity Model**
  - High fidelity
  - Each airport requires a different model
  - Data intensive

Weather/Capacity Simulation

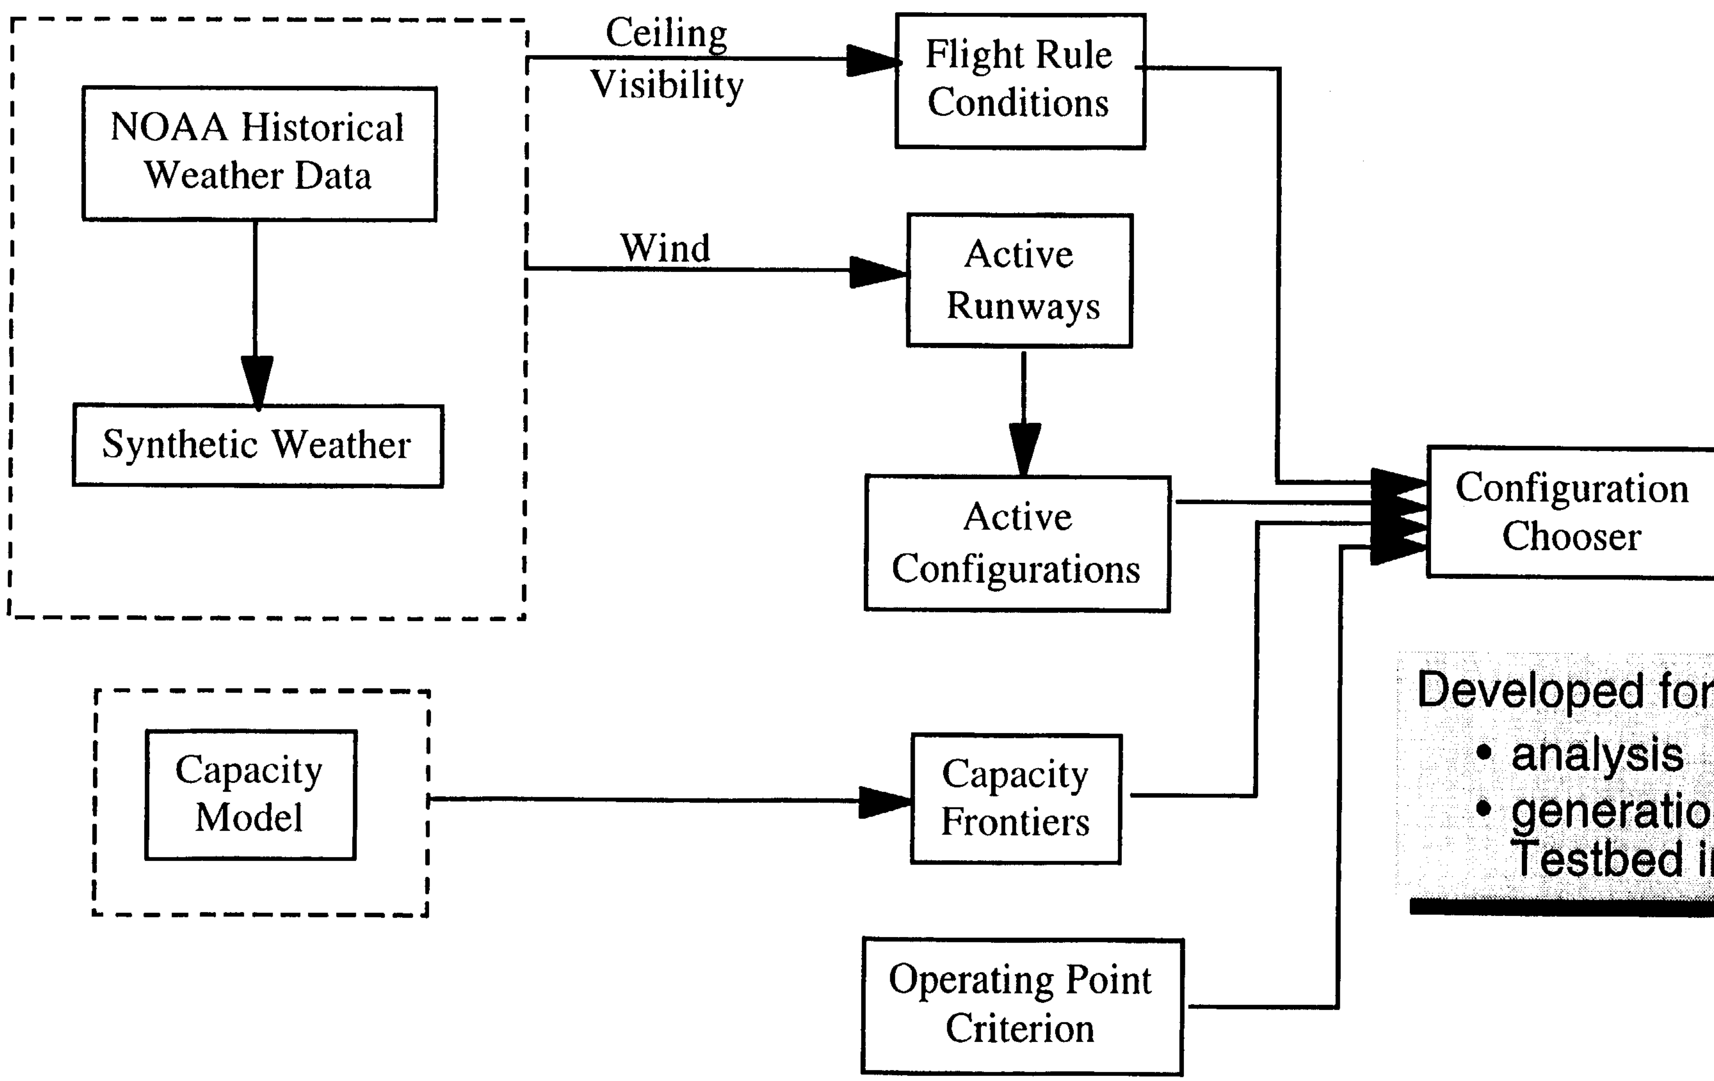

Developed for
- analysis
- generation of Testbed input

Example Coverage Results

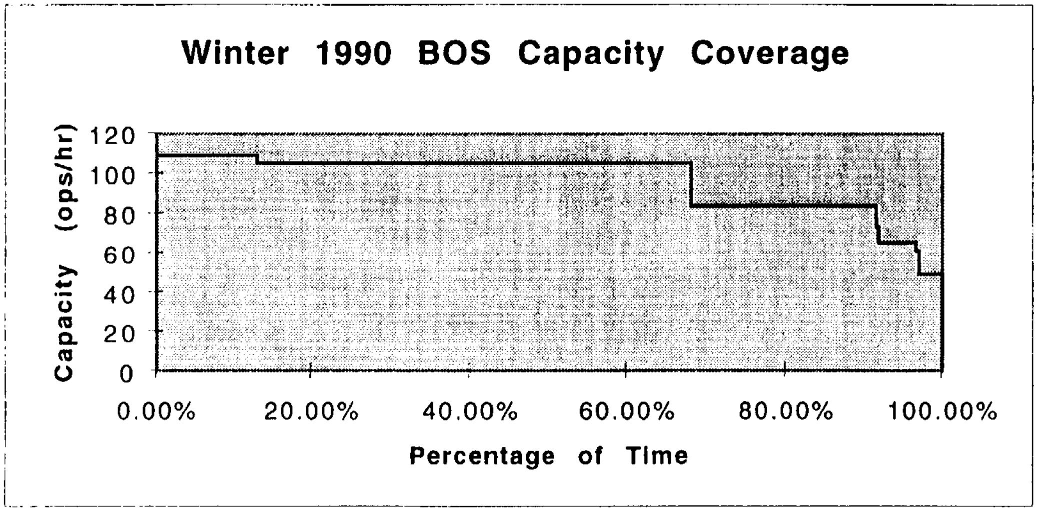

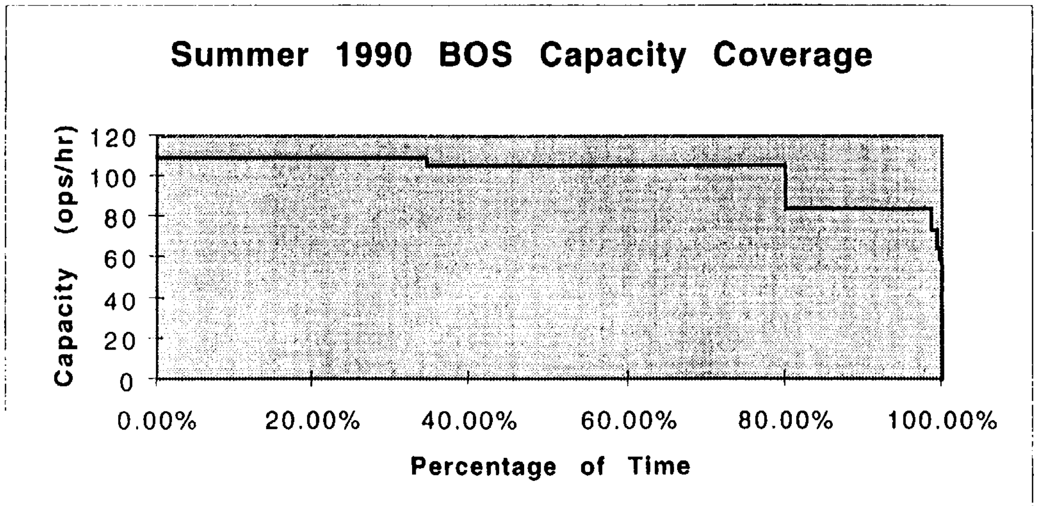

EPS model
50/50 arrival/departure mix

Example Distribution Results

BOS Capacity Probabilities

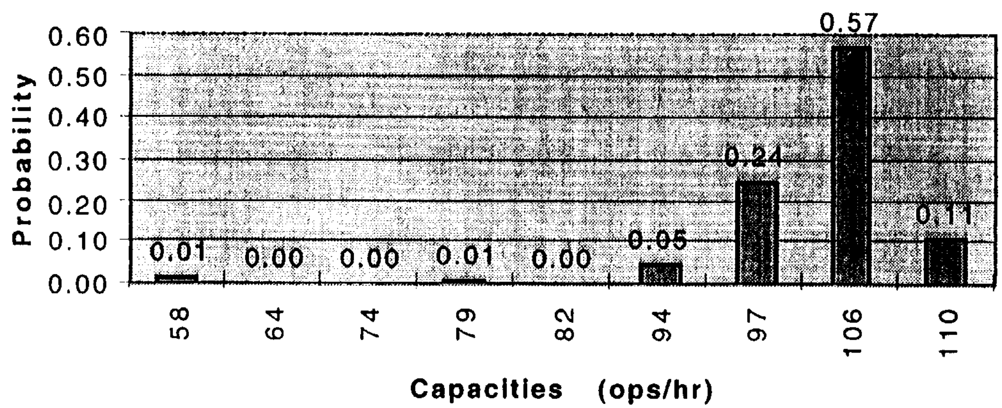

Single set of parallel runways

EPS
50/50 arrival/departure mix
Jan-Feb over 5 years

SFO Capacity Probabilities

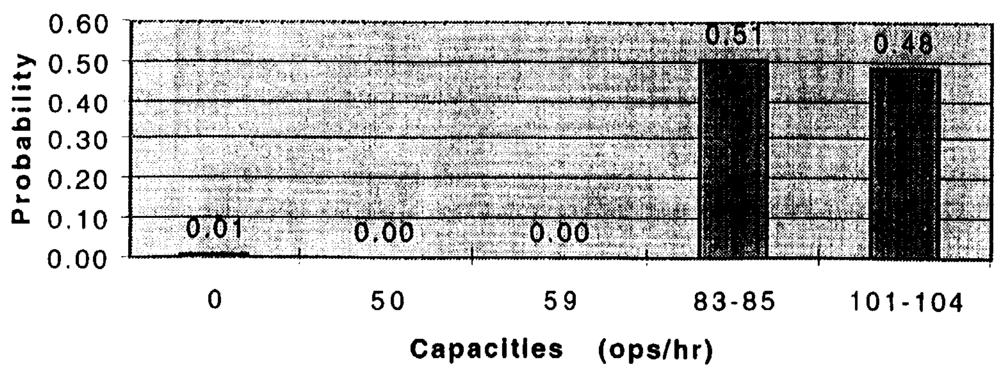

Multiple parallel runways

Planning, Simulation and Evaluation

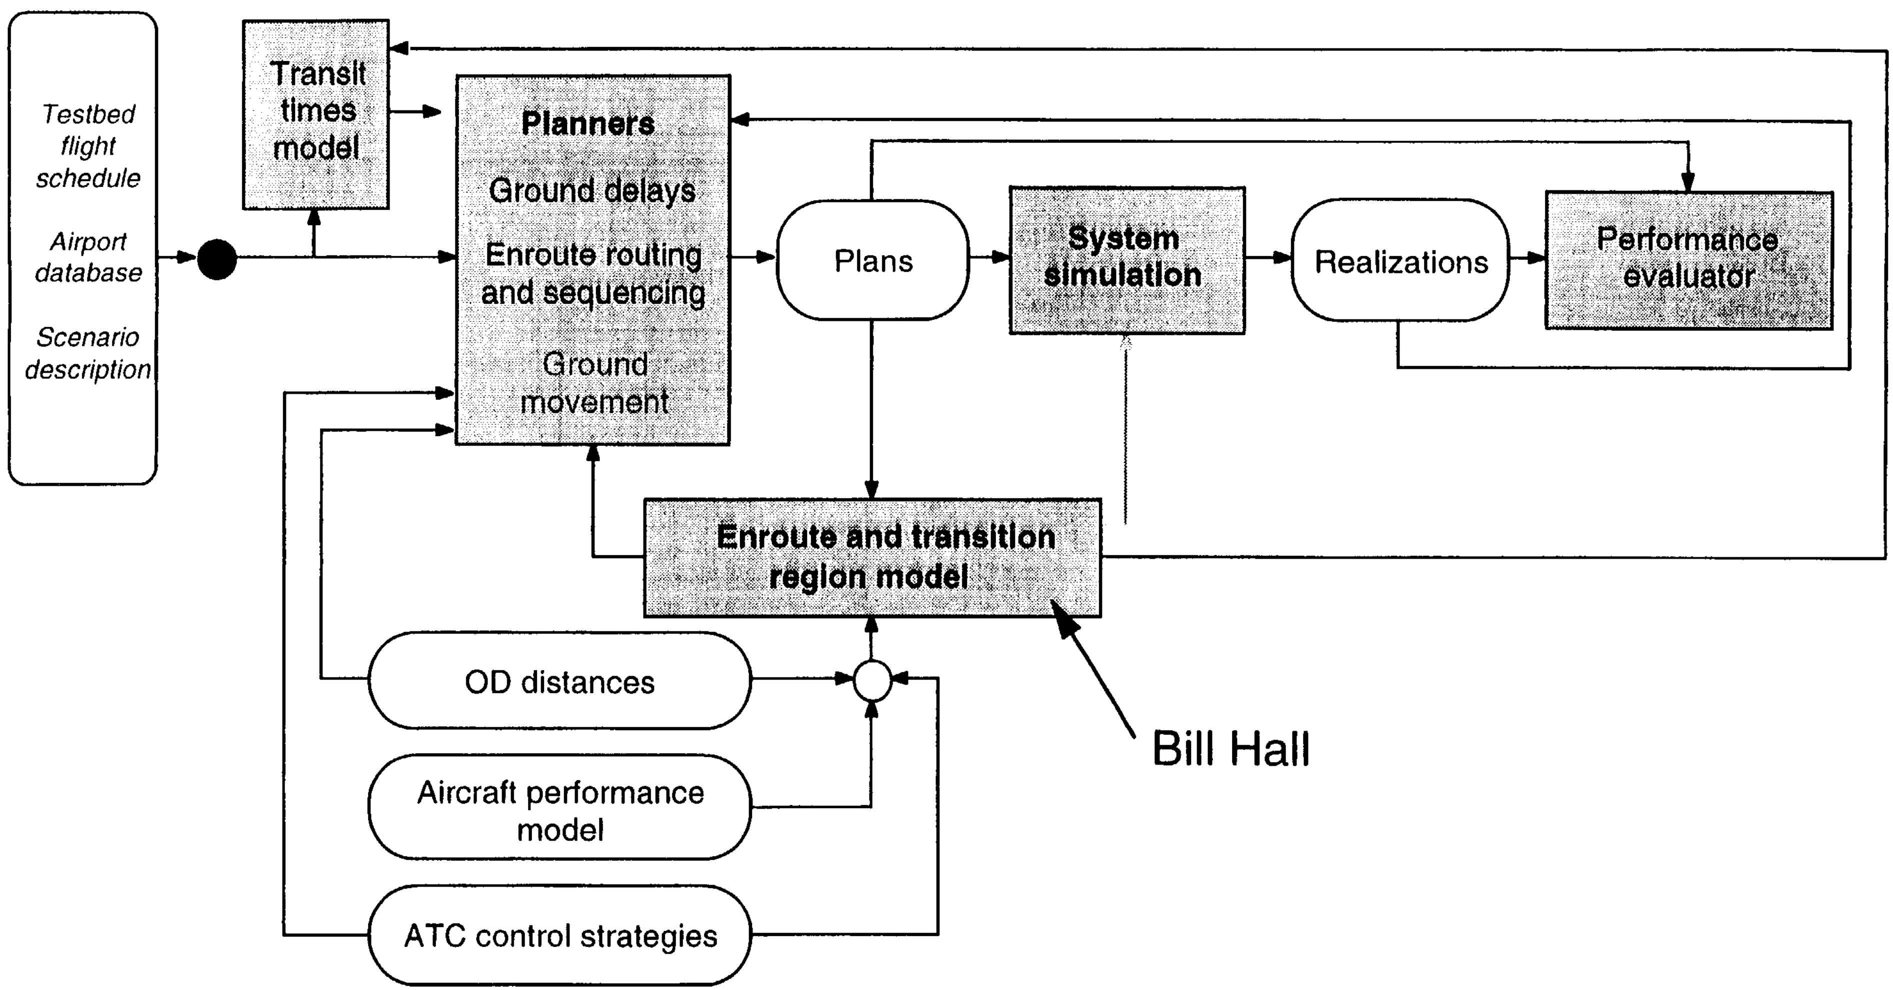

Tactical Enroute & Transition Area Model

DRAPER LABORATORY

	Allocation of Arrival Slots	Assignment of Arrival Slots to Individual Flights	Assignment of Departure Slots to Individual Flights	En Route Planning and Control	Transition Area, Terminal Area, Ground Movement Planning and Control
Centralized I	1a TFM system operator (FAA) allocates arrival slots to individual flights.	2a TFM system operator assigns slots; each airline may cancel and substitute flights	3a TFM system operator assigns departure slots to individual aircraft.	4a Airlines plan; TFM system operator controls.	5a Airlines plan; TFM system operator controls.
Partially Centralized (Current system)		2b Each airline suggests alternative assignment of the slots allocated to it (for its own flights only); TFM system operator approves or rejects.		4b Airlines plan; TFM system operator specifies a region in which airlines control their own aircraft; TFM system operator controls other regions and monitors globally for feasibility, conflicts.	
Partially Decentralized I	1b TFM system operator allocates sets of arrival slots to individual airlines.	2c Individual airlines allocate their own sets of slots among their own flights.	3b Airlines assigr slots to indivic TFM system o approves or re		
Partially Decentralized II		2d Airlines may trade slots among themselves. Each individual airline allocates its own set of slots among its flights.			
Decentralized I	1c TFM system operator informs airlines of the legal safe capacities at potentially congested airports.	2e Airlines may bargain among themselves for the legal safe capacities. They may cancel or delay flights, follow the original schedule, etc. within their "purchased" slots.	3c Airlines assign departure slots to individual aircraft.		5b Airlines plan; they also specify each aircraft's heading directly after departure; TFM system operator can approves or rejects heading; TFM system operator controls everything else.
Decentralized II	1d TFM system operator informs airlines about anticipated availability of capacities at congested airports.	2e Airlines decide what they will do. They may cancel or delay flights, follow the original schedule, etc.		4c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.	5c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.

Tactical Enroute & Transition Area Model Description

- **Higher fidelity enroute and transition area model**
- **Characteristics**
  - Air space congestion model
    - Measures
      - Interaircraft distances
      - Delay incurred due to avoidance maneuvers
      - Maneuvering workload
    - Enroute separation constraints enforced
  - Aircraft dynamics model
    - Velocity
    - Yaw and pitch rates
    - Pitch limits
  - Each aircraft is flown by a low fidelity pilot behavior model
    - Flies toward destination or away from problems as they occur/are anticipated
    - Does not plan ahead beyond current problem
  - Airport arrival rates and landing separation constraints
  - Transition region, terminal area separation constraints
  - Departure rates and separation constraints
  - Aircraft interaction outputs
    - Interaircraft distances
    - Delay incurred due to avoidance maneuvers
    - Maneuvers

Pilot Behavior Model

- **Local situation awareness: The pilot is assumed to have knowledge of the position and velocity of every aircraft in a local vicinity**
  - The projected trajectories of the aircraft are approximated by straight lines
  - In the pilot's moving frame of reference, his aircraft's trajectory is a point
  - The distance between the two aircraft as a function of time is the distance from the point representing the pilot's aircraft to the linear function representing the other aircraft's position
  - The interaircraft distance function is hyperbolic in time

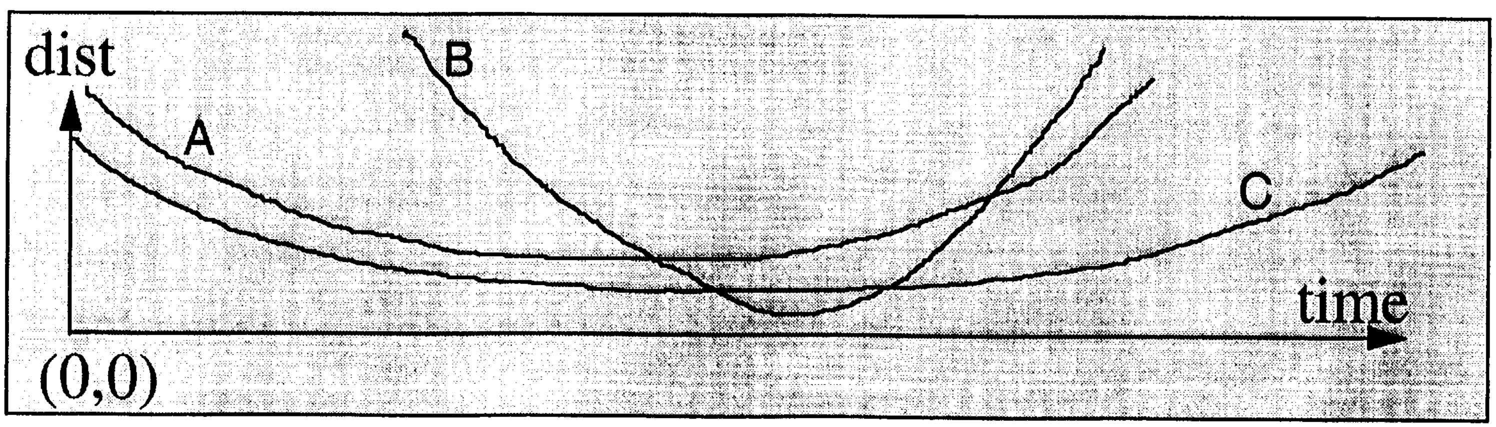

Interaircraft distances between ownship and aircraft A, B and C as a function of time

Importance Calculation

- The distance to the point on the hyperbola closest to (0,0) is the characteristic distance of the other aircraft
- The aircraft with the smallest characteristic distance is the most important (aircraft *C* below at time 0)
- If no aircraft's characteristic distance is below a predefined threshold, the pilot flies toward his destination

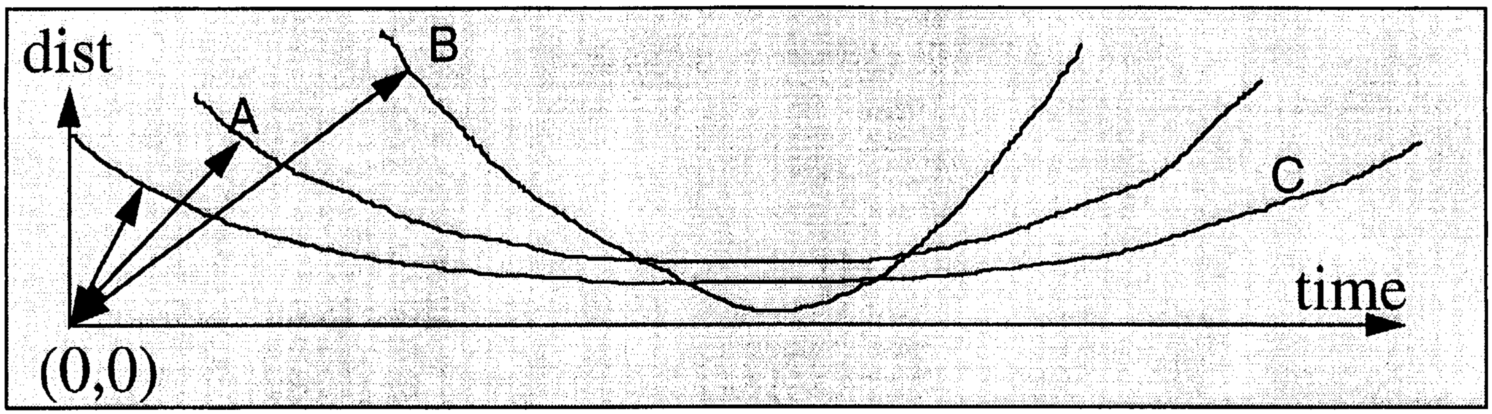

Avoidance actions act as a safety layer: real pilots and ATM system should prevent avoidance actions

Avoidance Action

- Pilot turns away from the position of the other aircraft at its point of closest approach in his frame of reference
- Pilot favors horizontal component of turn

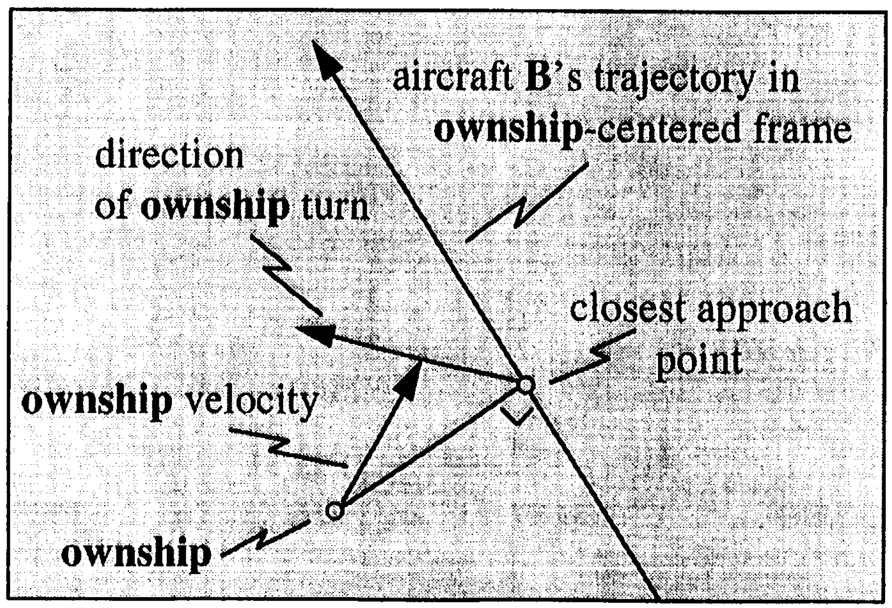

Enroute Stress Case Scenario

- 250 aircraft initialized within 100 km x 100 km x 10 km area to arrive simultaneously at center of region in five minutes

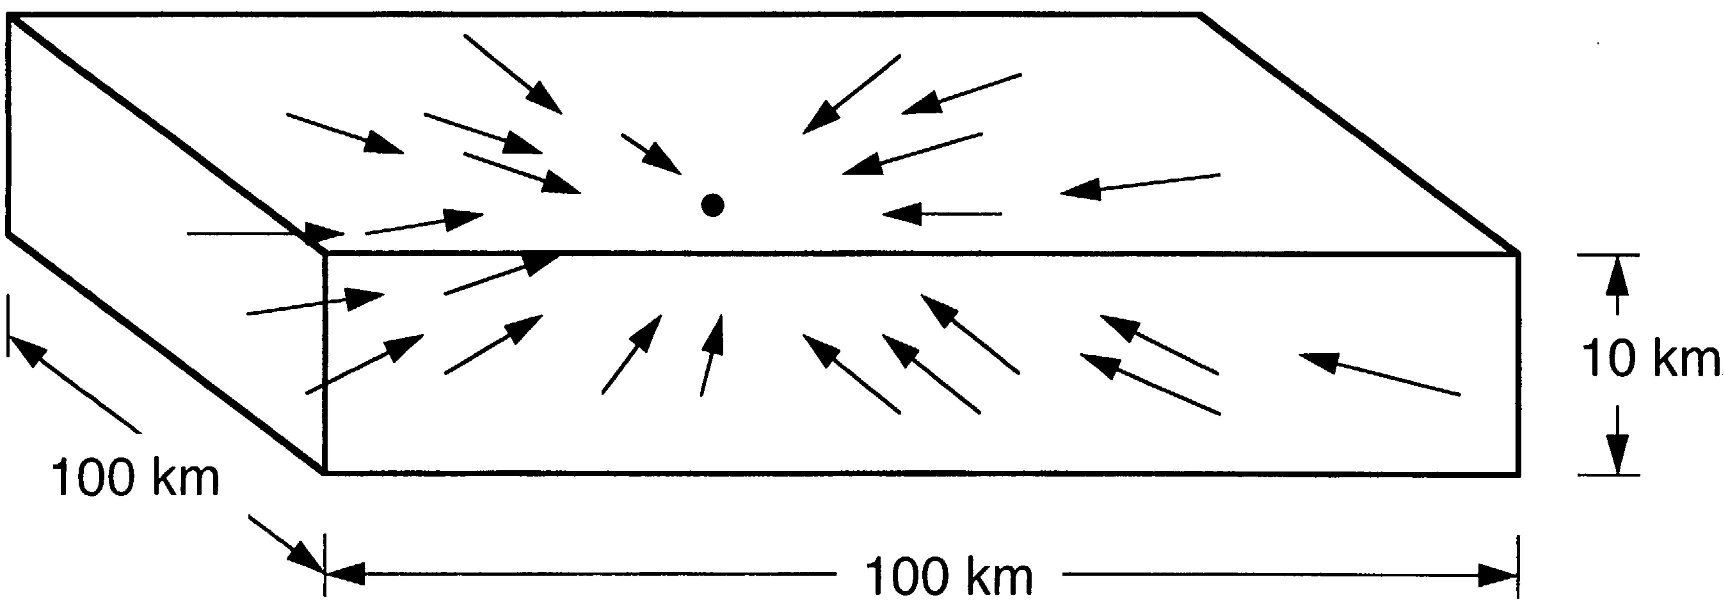

Enroute Stress Case Results

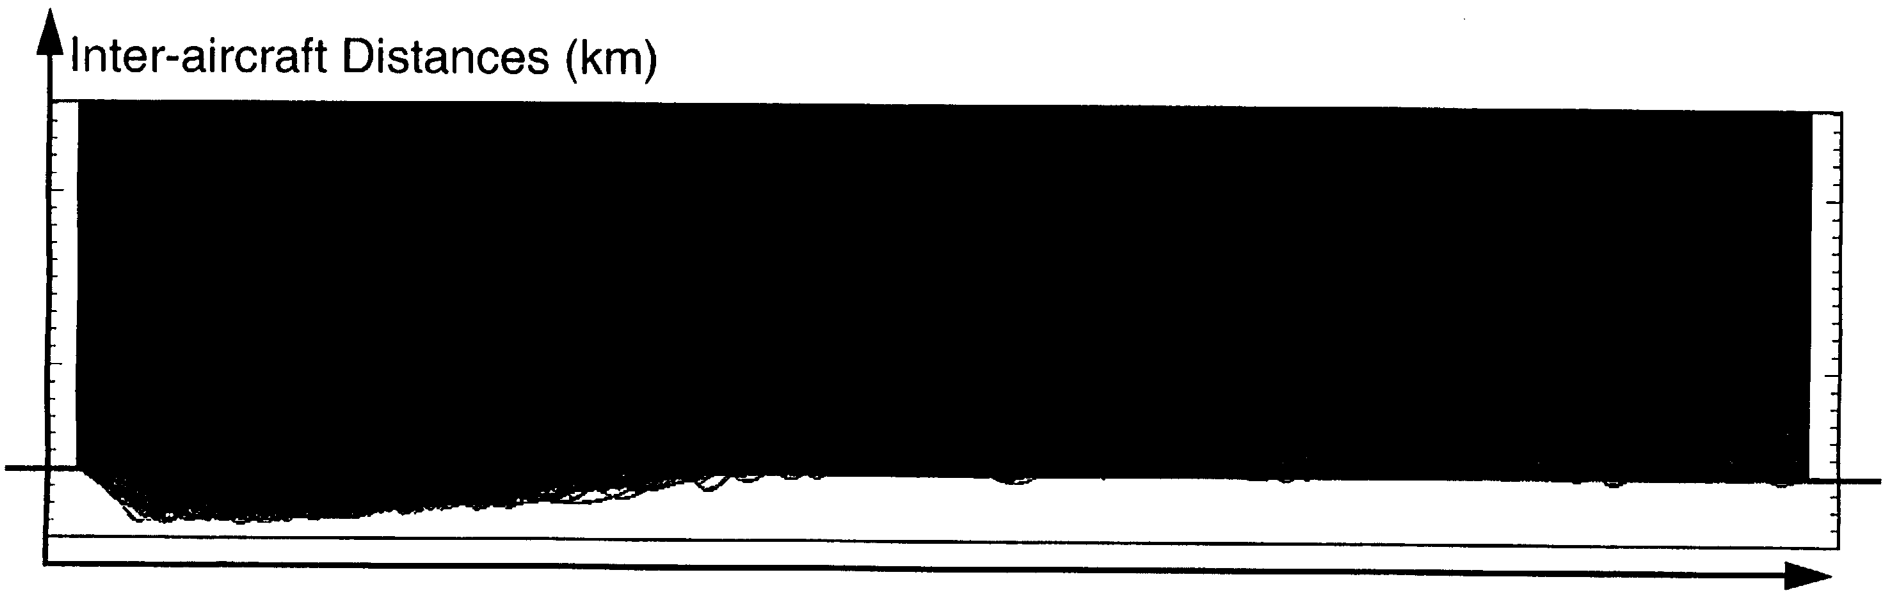

- **Run Statistics:**
  - Delay 133465 seconds (53%)
  - Maneuvers 54566 (one per 4.6 sec per a/c)

Enroute Example: Crossing Airways

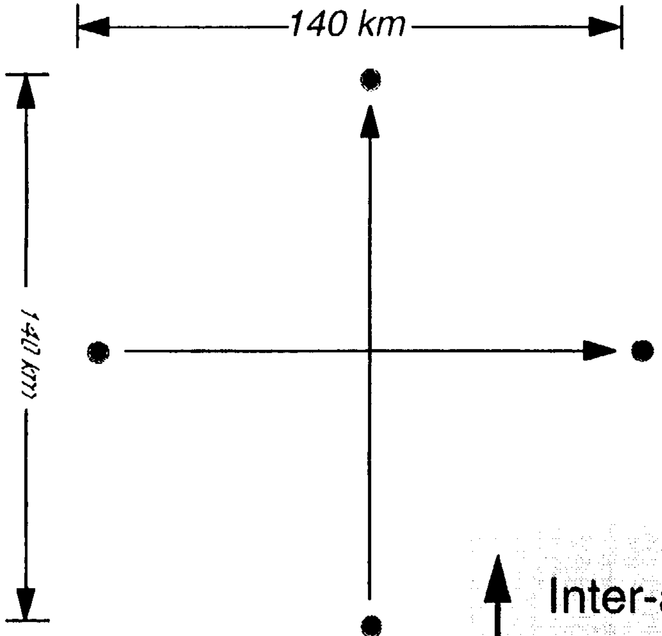

- Departures every 50 seconds for each airway
- Airways at same altitude
- Waypoints separate terminal area from enroute
- Delay is less than 1% per flight segment
- 4 km inter-aircraft distance constraint not violated

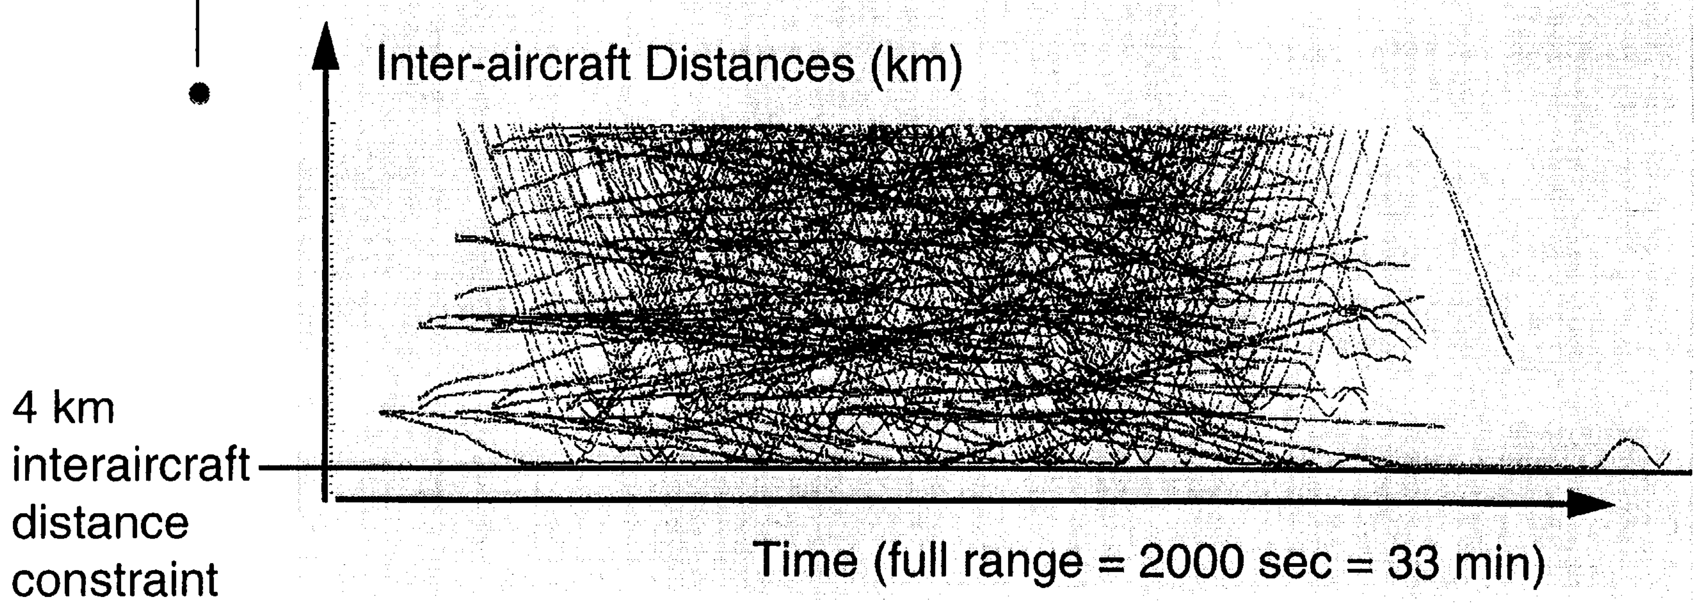

Coordination

- **Using distributed separation assurance alone**
  - Problem occurs when two or more aircraft try to attain the same waypoint at the same time (e.g., arrival feeder fix point into transition area)
  - By avoiding each other, neither gets to waypoint
  - If aircraft arrive sequentially, no problem
- **Solution: provide sequencing guidance information**
  - Do not have to force specific trajectories
    - e.g., free flight constrained by sequencing and spacing
  - Objective: correct spacing in time and distance from other aircraft when landing

Sequence Enforcement

- Assume a landing sequence has been determined
- Each aircraft gets an approach alignment waypoint (or controlled airspace handoff waypoint)

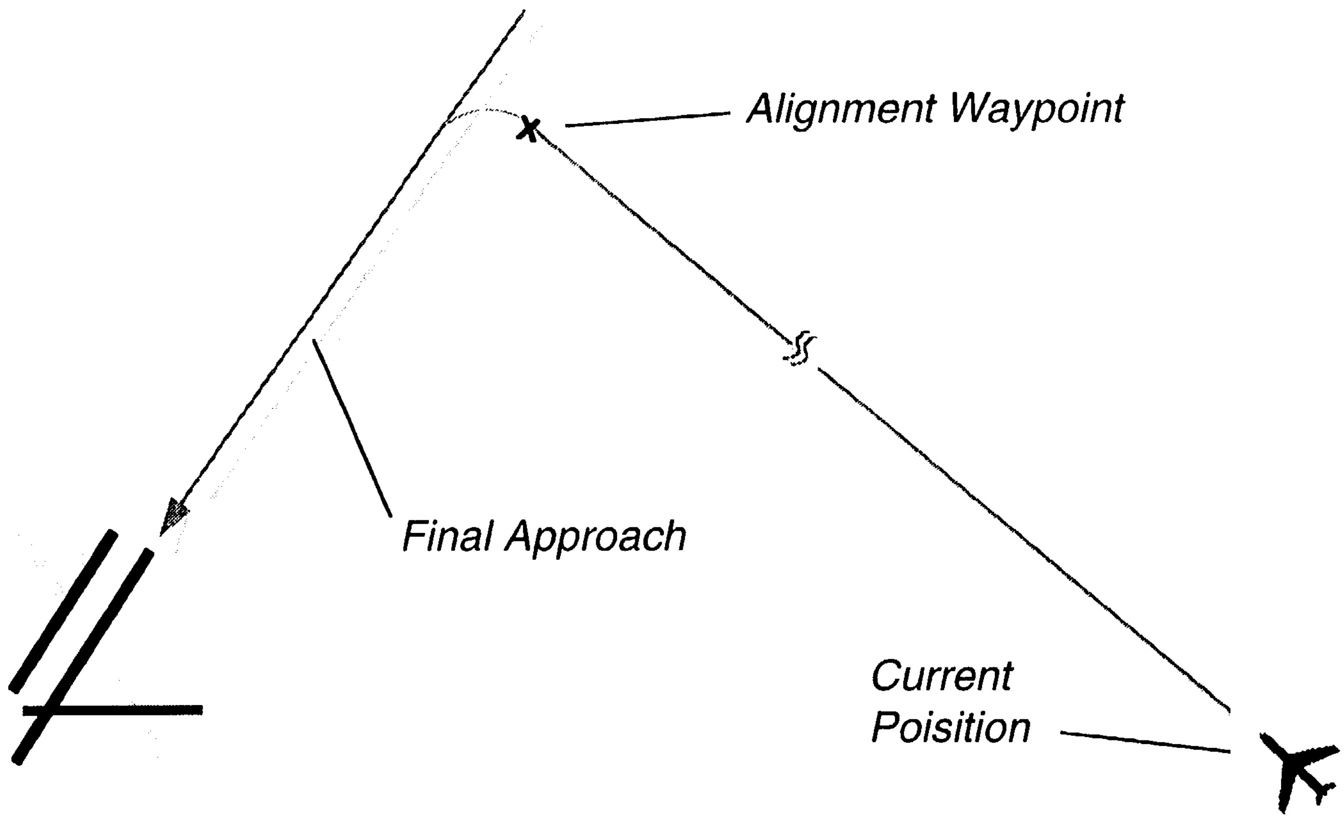

Sequence Enforcement

- Assign each aircraft a moving keep-out boundary
- Separation assurance mechanism enforces keep-out boundary

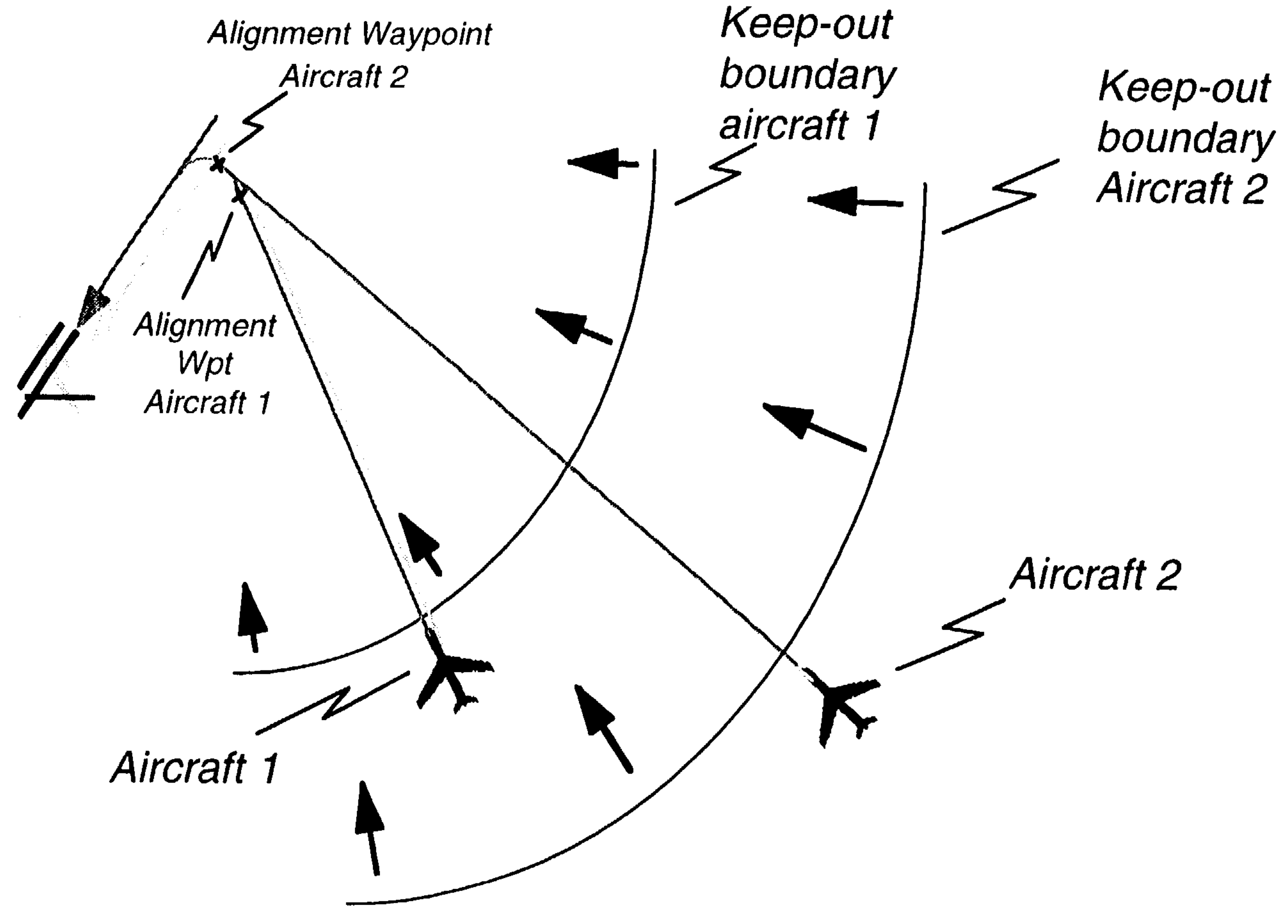

Future Directions

- **Sequencing will allow us to model effect on system performance of size of controlled regions**
  - Controlled region is where aircraft are under positive control
  - Controlled region defined by maximum keep-out radius set by system operator
  - Maximum keep-out radius expanded during peak hours for safety
  - During low demand periods, controlled region might contract to allow greater free flight flexibility

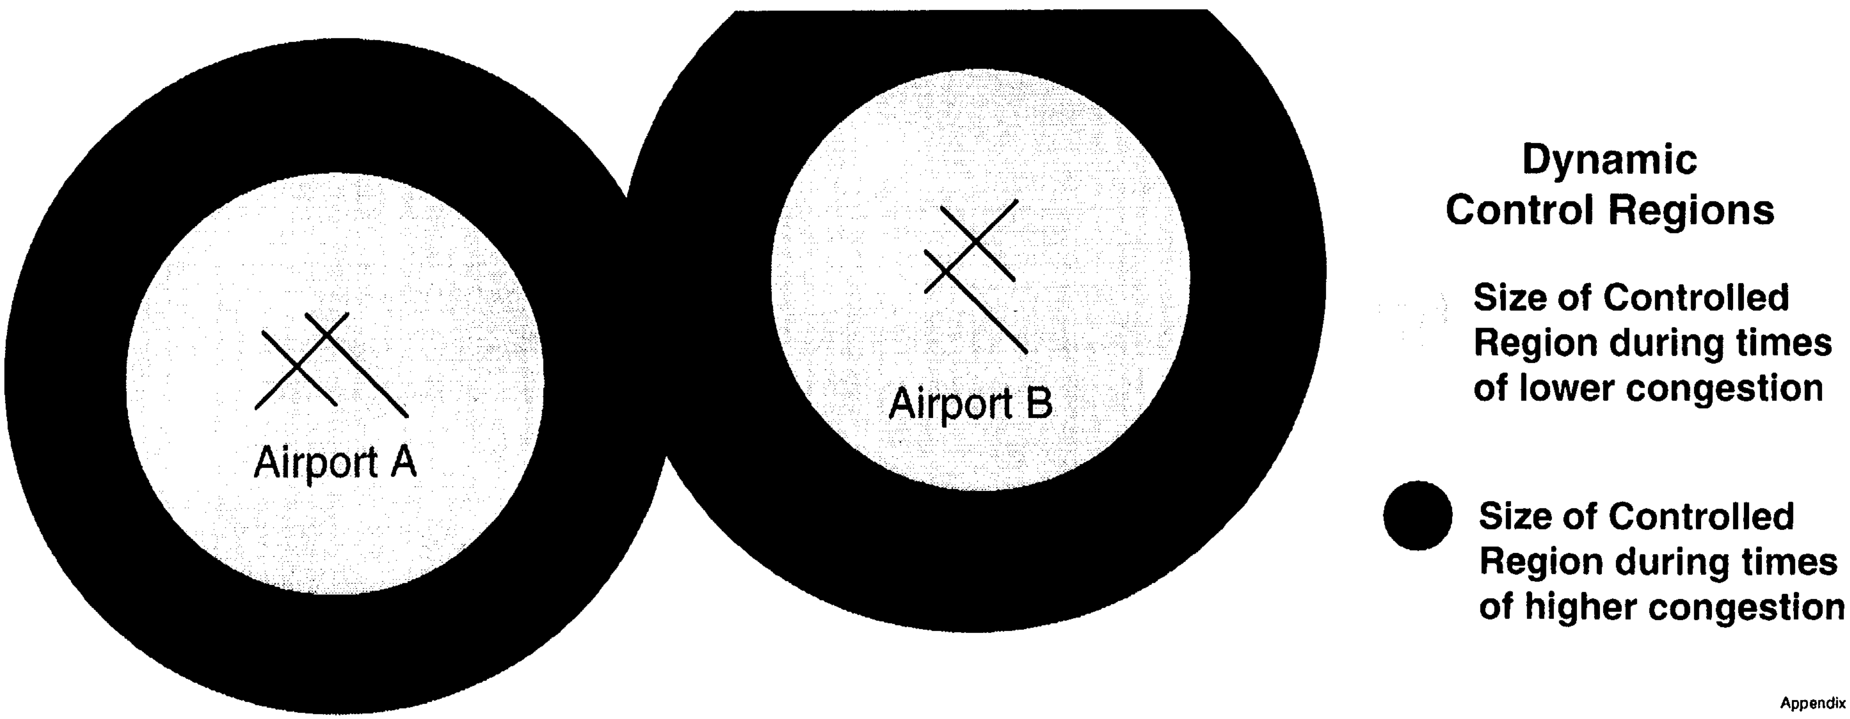

Future Directions (continued)

- **Combine higher fidelity enroute and transition area model with weather models to explore enroute congestion problems**
  - Create dynamic controlled regions in enroute congestion areas
- **Add realistic user-preferred routing**
  - Take into account winds aloft, weather, and congested areas
- **Cost and safety analysis**

Planning, Simulation and Evaluation

Prof. Amedeo Odoni
Mark Abramson
Paul Carlson
Sarah Stock

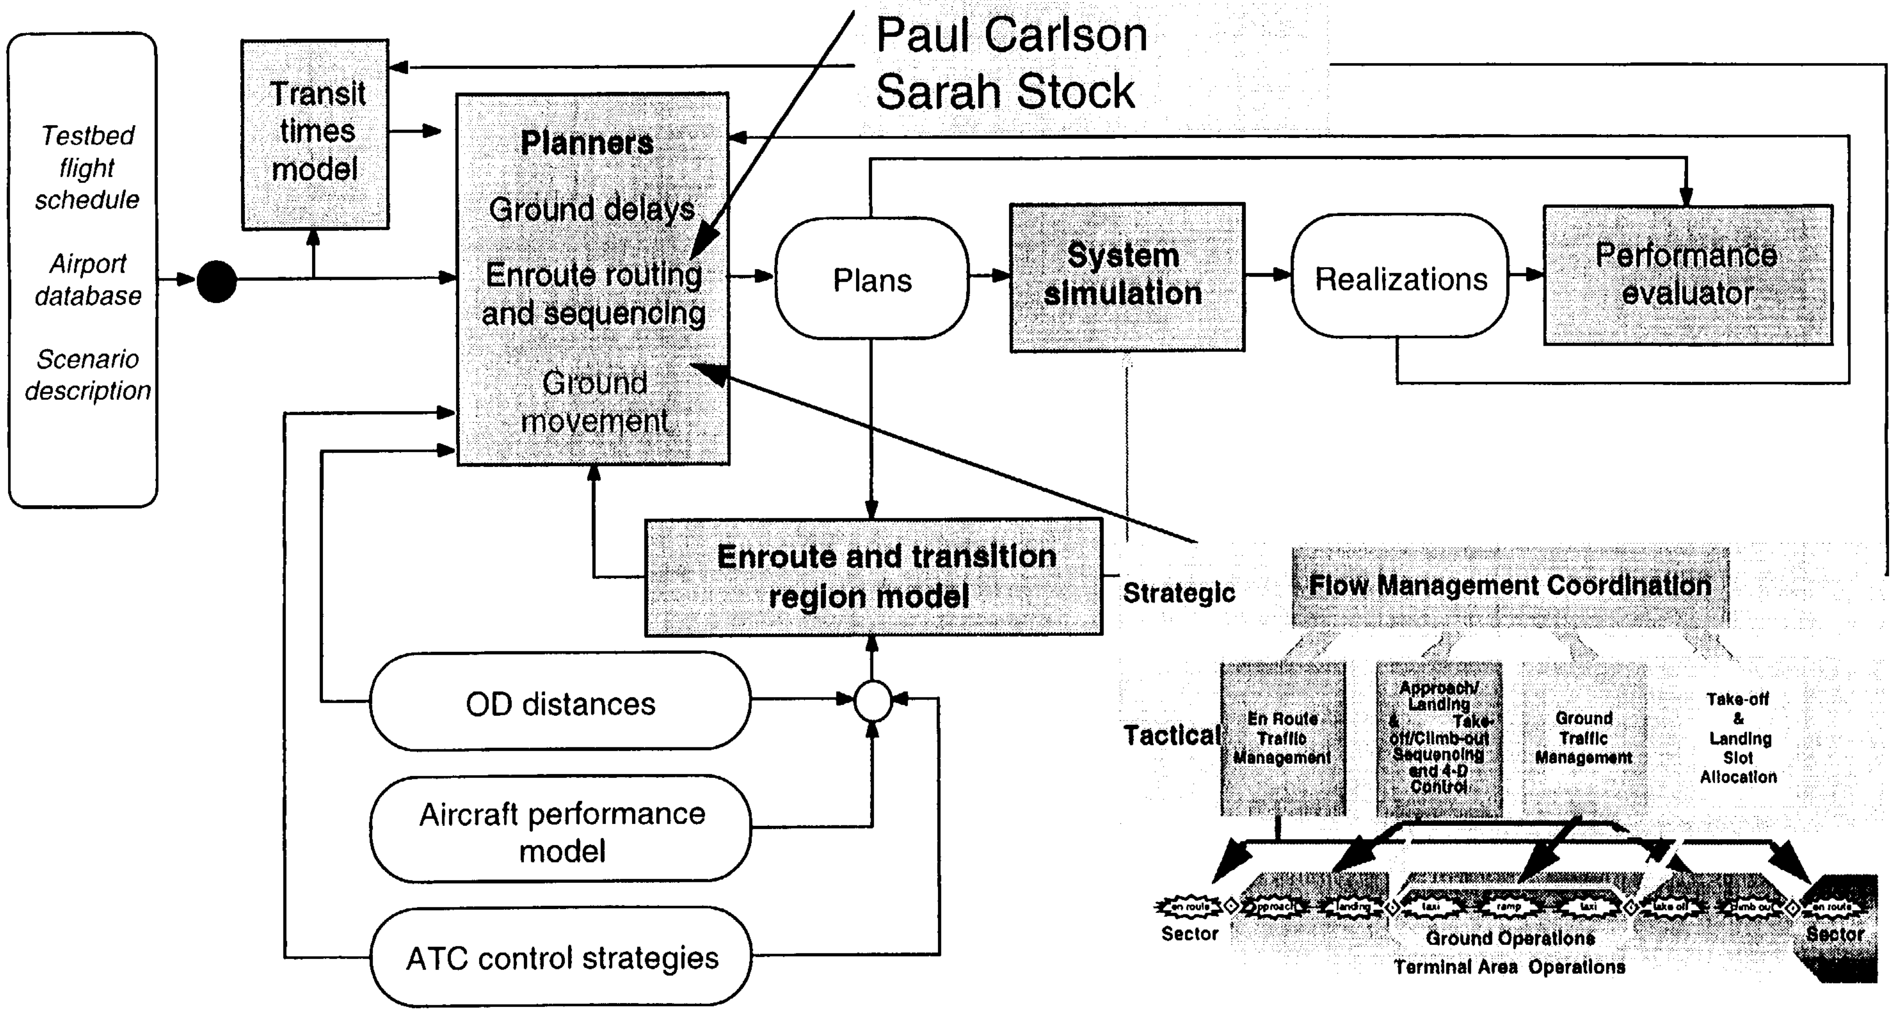

Arrival Slot Scheduling Heuristic

	Allocation of Arrival Slots	Assignment of Arrival Slots to Individual Flights	Assignment of Departure Slots to Individual Flights	En Route Planning and Control	Transition Area, Terminal Area, Ground Movement Planning and Control
Centralized I	1a TFM system operator (FAA) allocates arrival slots to individual flights.	2a TFM system operator assigns slots; each airline may cancel and substitute flights	3a TFM system operator assigns departure slots to individual aircraft.	4a Airlines plan; TFM system operator controls.	
Partially Centralized (Current system)		2b Each airline suggests alternative assignment of the slots allocated to it (for its own flights only); TFM system operator approves or rejects.			
Pa De	1b TFM system operator allocates sets of arrival slots to individual airlines.	2c Individual airlines allocate their own sets of slots among their own flights.	3b Airlines assign departure slots to individual aircraft; TFM system operator approves or rejects.	4b Airlines plan; TFM system ...ies a region ...es control ...aft; TFM ...r controls ...nd ...lly for feasibility, conflicts.	5a Airlines plan; TFM system operator controls.
Partially Decentralized II		2d Airlines may trade slots among themselves. Each individual airline allocates its own set of slots among its flights.			
Decentralized I	1c TFM system operator informs airlines of the legal safe capacities at potentially congested airports.	2e Airlines may bargain among themselves for the legal safe capacities. They may cancel or delay flights, follow the original schedule, etc. within their "purchased" slots.	3c Airlines assign departure slots to individual aircraft.		5b Airlines plan; they also specify each aircraft's heading directly after departure; TFM system operator can approves or rejects heading; TFM system operator controls everything else.
Decentralized II	1d TFM system operator informs airlines about anticipated availability of capacities at congested airports.	2e Airlines decide what they will do. They may cancel or delay flights, follow the original schedule, etc.		4c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.	5c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.

Background

- **Andreatta-Brunetta-Guastalla (ABG) heuristic algorithm developed at University of Padua**
  - Presented as a (strategic) ground holding algorithm
  - Flights are assigned landing slots based on a flight's ***priority***
    - Does the flight have connections?
    - How much delay has already been assigned?
  - Near optimal performance achieved with greatly reduced (relative to optimal Bertsimas-Stock algorithm) processing time in preliminary test cases
- **Guastalla continuing research on ABG at MIT**
  - Used to develop an initial solution for optimal ground holding algorithm to decrease algorithm CPU time
- **Draper integrating ABG into the simulation testbed**
  - Strategic use as a ground holding algorithm to demonstrate effect of different priority setting schemes
  - Tactical use in arrival slot scheduling

Sample Heuristic Performance Results

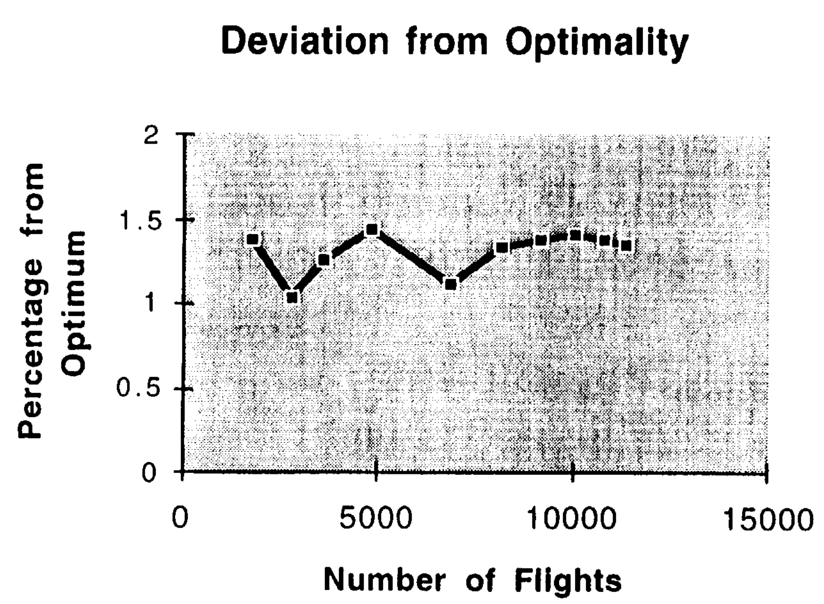

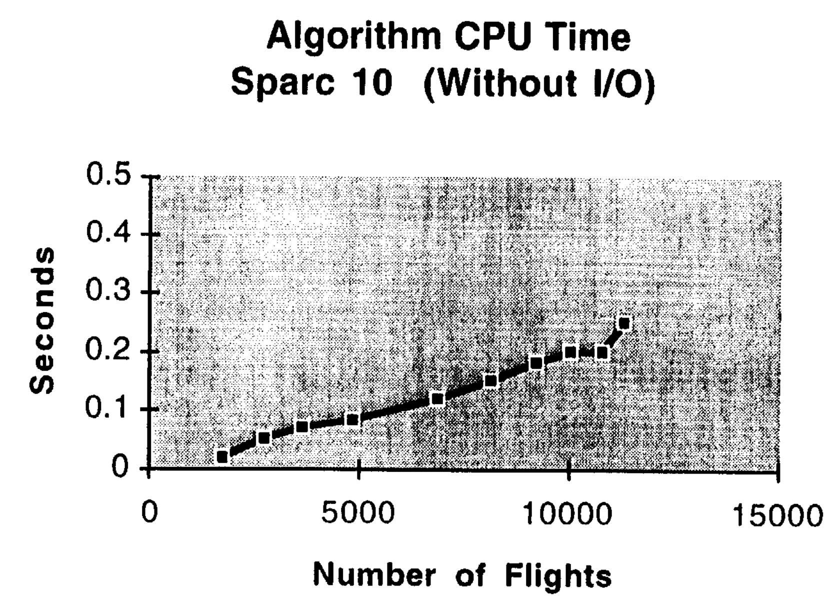

Caveats:

- Synthetic schedules
- Simple cost functions

Conclusion:

- Excellent results indicate worth pursuing as a tactical algorithm

ABG Flowchart

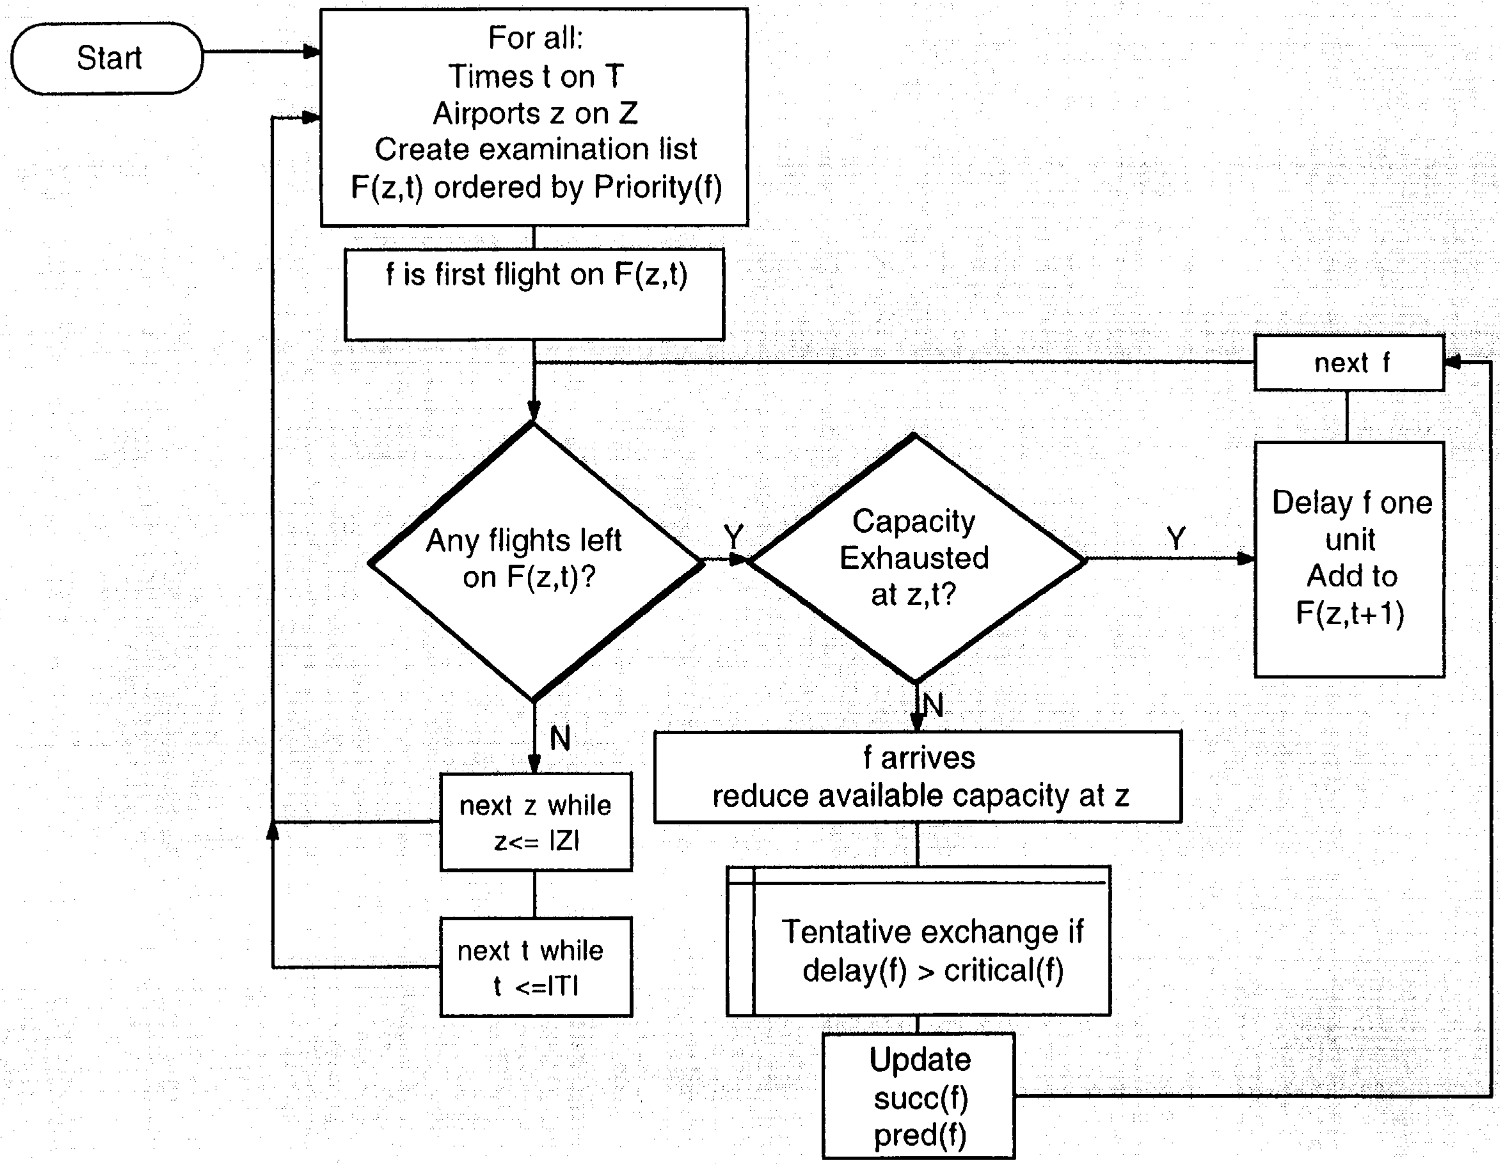

Experiments

- **Test strategic planning capability of ABG**
- **Specify priority schedules emphasizing example airline concerns**
  - Aircraft size: higher passenger capacity flights get preference
  - Preference to flights with connections
  - Connections + Delay already assigned
  - Size + Connections + Delay
- **Specify new metrics for measuring improvements brought by ABG over non-prioritized strategic plan (e.g., FCFS)**
- **Closer relationship with airlines required to identify useful priorities measures, schedules and costs**

Extensions : Tactical Use

- **Fast**
  - Use, as needed, during the day as the forecast of future system demand and capacity changes
- **Flexible**
  - Decisions are made on the basis of priorities—can model a range of airline operational policies
- **Provides good plans**
  - Within a few percent of the optimal solution in synthetic scenarios

Extensions: Tactical Use (Continued)

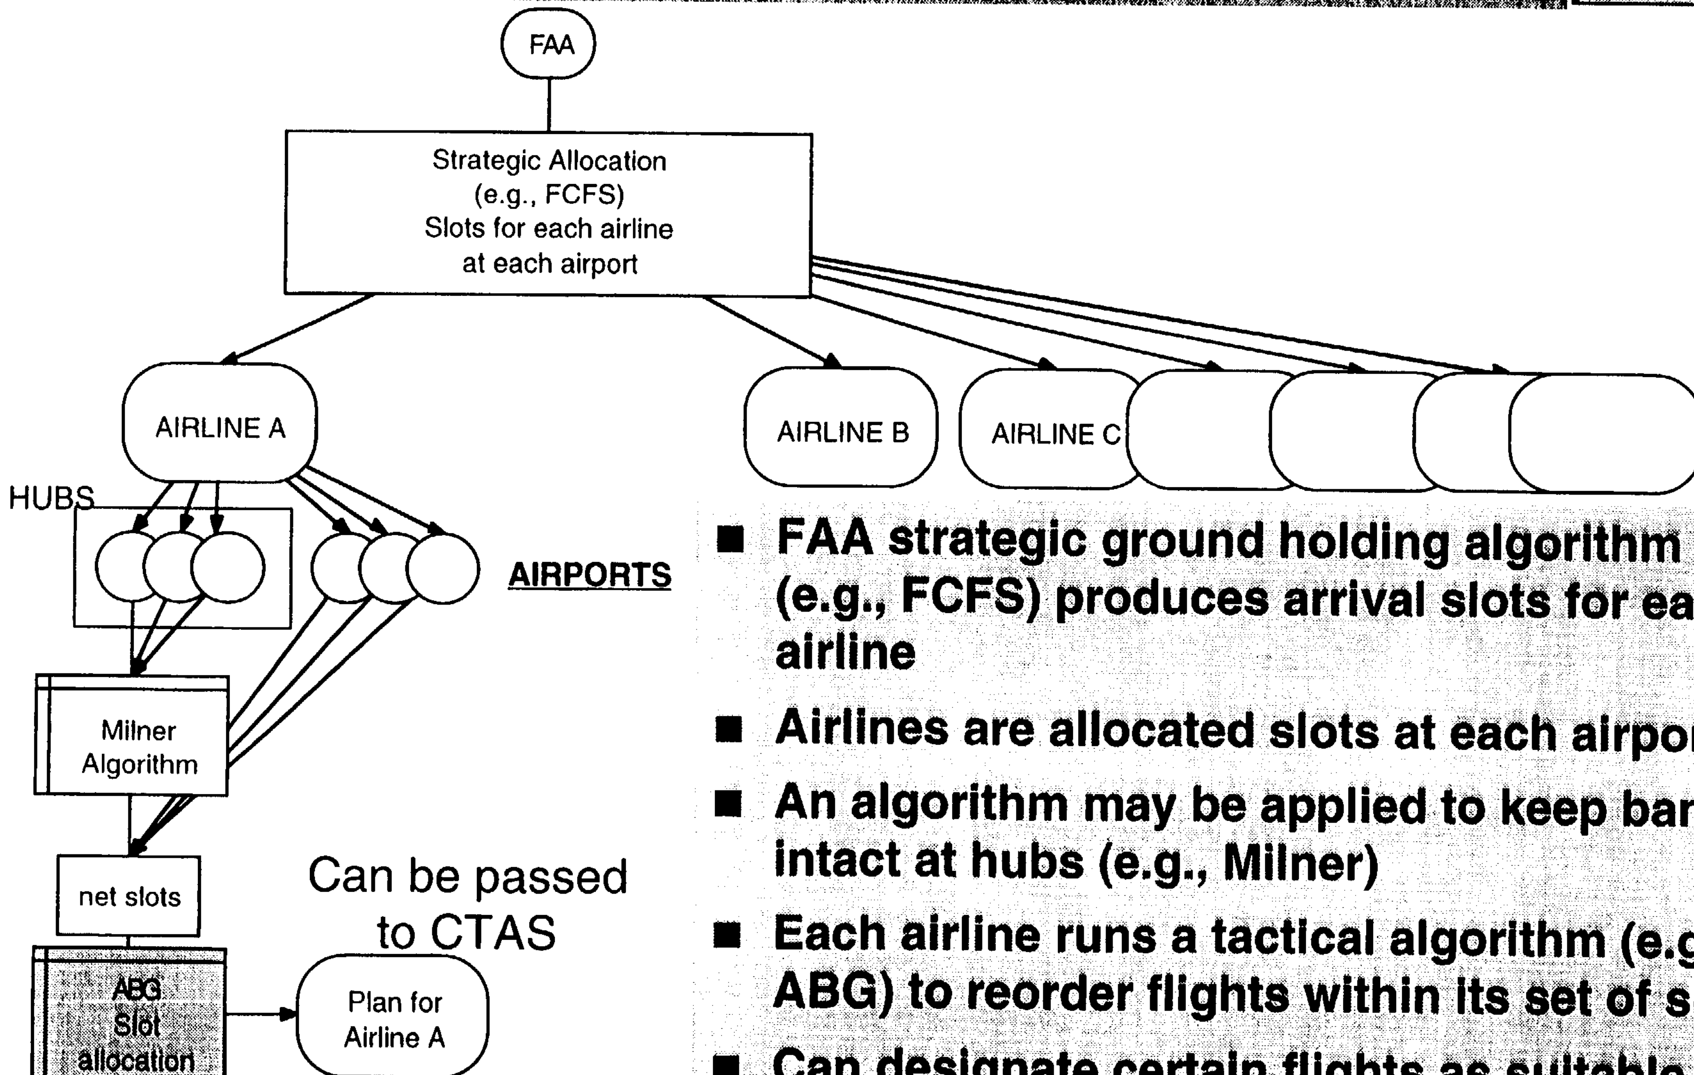

- FAA strategic ground holding algorithm (e.g., FCFS) produces arrival slots for each airline
- Airlines are allocated slots at each airport
- An algorithm may be applied to keep banks intact at hubs (e.g., Milner)
- Each airline runs a tactical algorithm (e.g., ABG) to reorder flights within its set of slots
- Can designate certain flights as suitable for cancellation

Other Extensions

- **User preferred routing**
  - Reduced transit times give ABG a window of time to schedule a flight without incurring any tardiness
- **Priorities as a function of tardiness amount**
- **Priorities that model hub operations, i.e., keeping banks together**

Airline Tactical Decision-Making for Slot Utilization

DRAPER LABORATORY

	Allocation of Arrival Slots	Assignment of Arrival Slots to Individual Flights	Assignment of Departure Slots to Individual Flights	En Route Planning and Control	Transition Area, Terminal Area, Ground Movement Planning and Control
Centralized I	1a TFM system operator (FAA) allocates arrival slots to individual flights.	2a TFM system operator assigns slots; each airline may cancel and substitute flights	3a TFM system operator assigns departure slots to individual aircraft.	4a Airlines plan; TFM system operator controls.	5a Airlines plan; TFM system operator controls.
Partially Centralized (Current system)		2b Each airline suggests alternative assignment of the slots allocated to it (for its own flights only); TFM system operator approves or rejects.		4b Airlines plan; TFM system operator specifies a region in which airlines control their own aircraft; TFM system operator controls other regions and monitors globally for ...y, conflicts.	
Partially Decentralized I	1b TFM syst... allocates s... slots to in... airlines.	2c Individual airlines allocate their own sets of slots among their own flights.	3b Airlines assign departure slots to individual aircraft; TFM system operator approves or rejects.		
Partially Decentralized II		... among themselves. Each individual airline allocates its own set of slots among its flights.			
Decentralized I	1c TFM system operator informs airlines of the legal safe capacities at potentially congested airports.	2e Airlines may bargain among themselves for the legal safe capacities. They may cancel or delay flights, follow the original schedule, etc. within their "purchased" slots.	3c Airlines assign departure slots to individual aircraft.		5b Airlines plan; they also specify each aircraft's heading directly after departure; TFM system operator can approves or rejects heading; TFM system operator controls everything else.
Decentralized II	1d TFM system operator informs airlines about anticipated availability of capacities at congested airports.	2e Airlines decide what they will do. They may cancel or delay flights, follow the original schedule, etc.		4c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.	5c Airlines plan and control their own aircraft; TFM system operator monitors for feasibility, conflicts.

Airline Tactical Decision-Making for Slot Utilization II

- **Develop optimization models for airline real-time slot assignment**
- **Two principal models:**
  - Each flight treated independently ("IF model")
  - Interdependencies/connections among flights taken into account ("C/D model")
- **Points to substantial differences between the two cases, with regard to flight delay strategies and flight cancellations**
- **Suggests significant benefits from partial decentralization of TFM for airlines with large numbers of inter-connecting flights and hub operations**
- **Milner (1995) Dynamic Slot Allocation with Airline Participation, PhD Thesis, 1995; Carlson (1996)**

Model Formulation

NOTATION

Input variables

I a set of flights, $i \in I$

B a set of banks, $b \in B$

T a set of discrete time periods, $t \in T$

M_t the number of slots available in period t

w_{bt} the cost of completing bank b at time t

c_{i0} the cost of canceling flight i

d^s_{it} the reward received for rescheduling separated flight i at time t

Decision Variables

z_{bt} 1 if bank b is assigned to be completed by time t; 0 otherwise

y_{it} 1 if flight i is assigned to a slot by time t; 0 otherwise

y_{i0} 1 if flight i is either outright canceled or is separated from its bank; 0 otherwise

y^s_{it} 1 if flight i is separated from its bank and is assigned a slot by time t; 0 otherwise

MODEL

Minimize $$\sum_b \sum_t w_{bt}(z_{bt} - z_{bt-1}) + \sum_i c_{i0} y_{i0} - \sum_i \sum_t d^s_{it}(y^s_{it} - y^s_{it-1}) \qquad (4.3a)^1$$

rewritten as

$$\sum_{i=1}^{I} \sum_{t=1}^{T-1} y^s_{it}(d^s_{it+1} - d^s_{it}) + \sum_{i=1}^{I} y^s_{iT}(-d^s_{iT}) + \sum_{i=1}^{I} y_{i0} c_{i0} + \sum_{b=1}^{B} \sum_{t=1}^{T-1} z_{bt}(w_{bt} - w_{bt+1}) + \sum_{b=1}^{B} z_{bT} w_{bT}$$

Model Formulation (continued)

subject to

$$\sum_i \left(y_{it} - y_{it+1}\right) + \left(y^s_{it} - y^s_{it+1}\right) \leq M_t \qquad \forall\, t \qquad (4.3b)$$

$$y_{it} \geq y_{it+1} \qquad \forall\, i, t \qquad (4.3c)$$

$$y^s_{it} \geq y^s_{it+1} \qquad \forall\, i, t \qquad (4.3d)$$

$$z_{bt} \leq y_{it} + y_{i0} \qquad \forall\, i \in b, \forall\, b, t \qquad (4.3e)^2$$

$$z_{bt} \geq z_{bt+1} \qquad \forall\, b, t \qquad (4.3f)$$

$$y_{it} \leq z_{b't} + y_{i0} \qquad \forall\, b' < b, \forall\, t, \forall\, i \in b, \forall\, b \qquad (4.3g)^3$$

$$z_{bT} = 1 \qquad \forall\, b \qquad (4.3h)$$

$$y^s_{it} \leq y_{i0} \qquad \forall\, i, t \qquad (4.3i)$$

$$y_{it}, y^s_{it}, y_{i0}, z_{bt} \in \{0,1\} \qquad (4.3j)^4$$

Arrival Pattern Used in Milner's Experiments

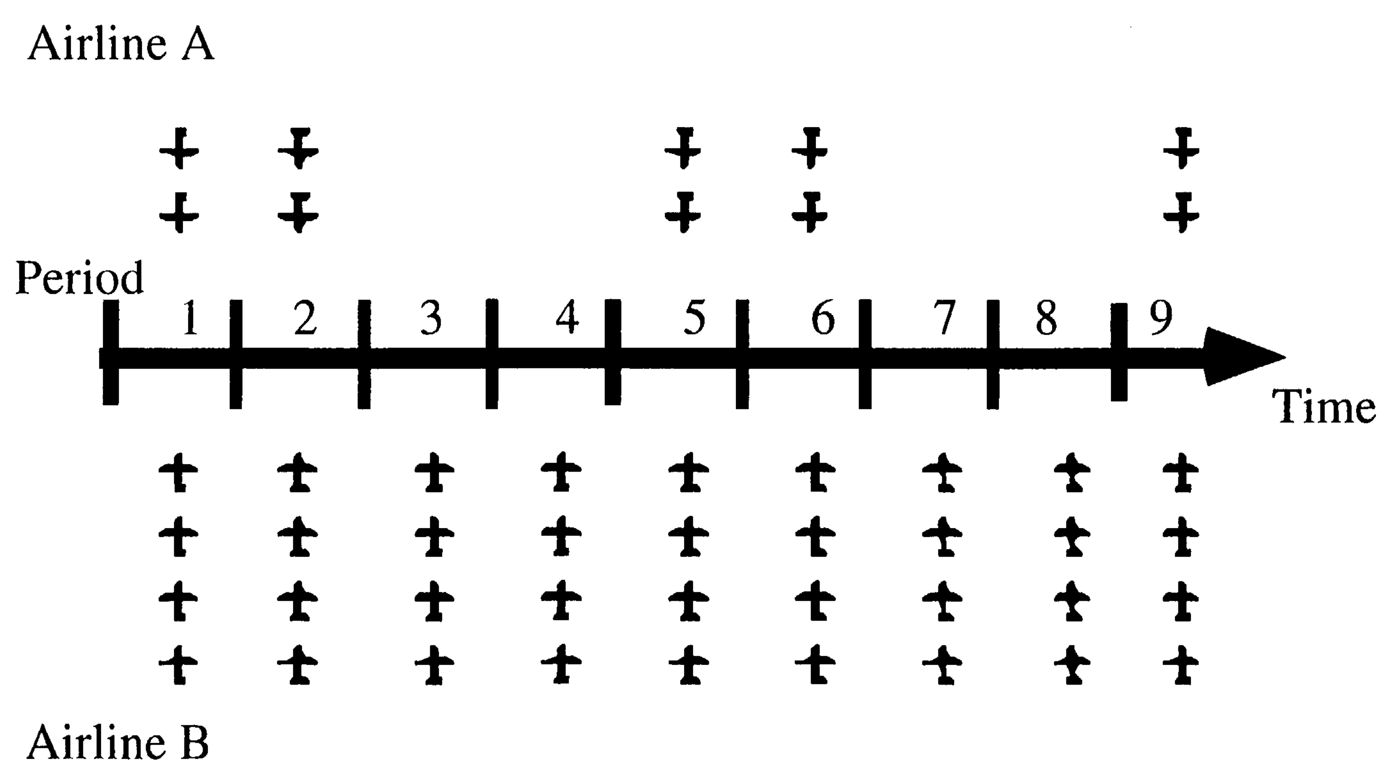

Canceled Flights vs. Delay Cost

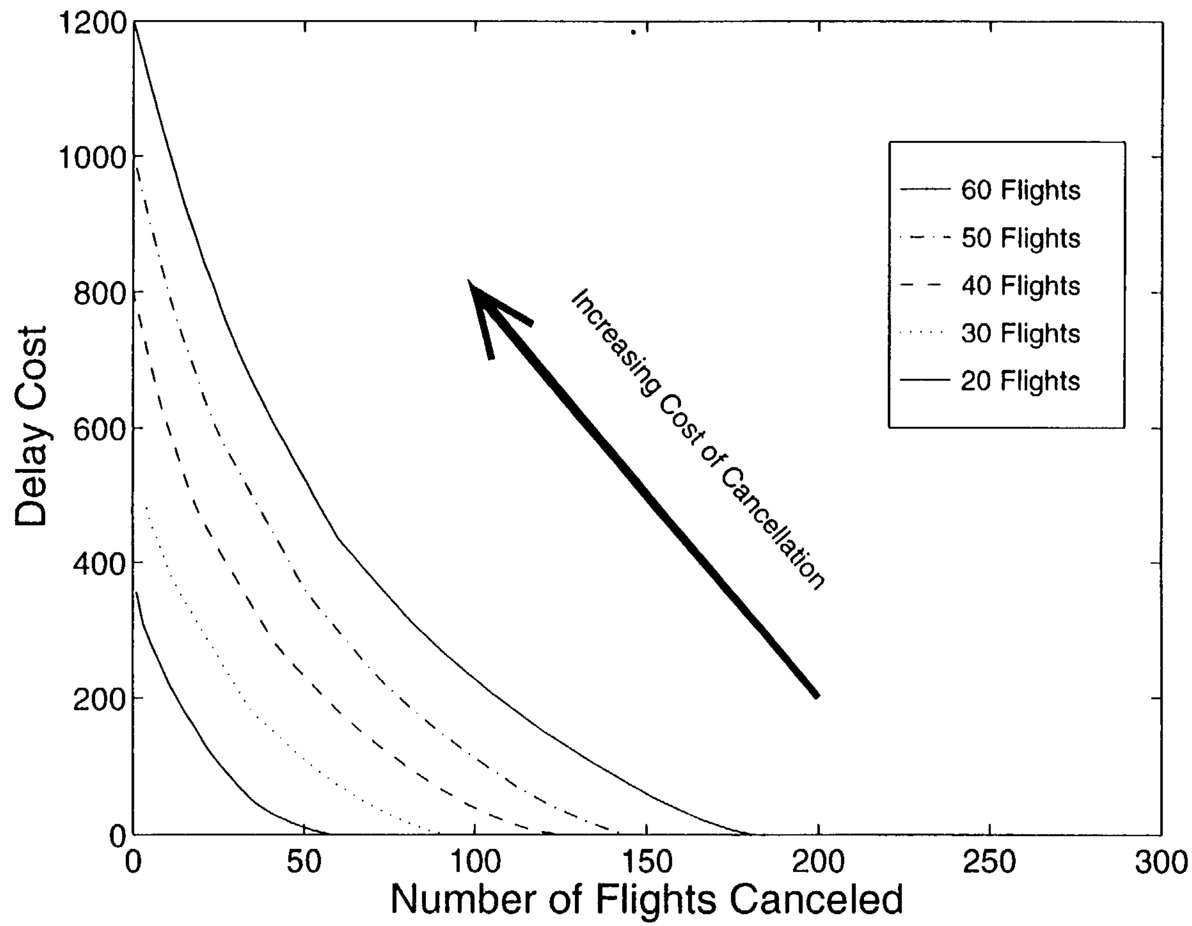

Number of canceled flights as a function of delay cost, with changing cost of cancellation

Comparison of IF and C/D models

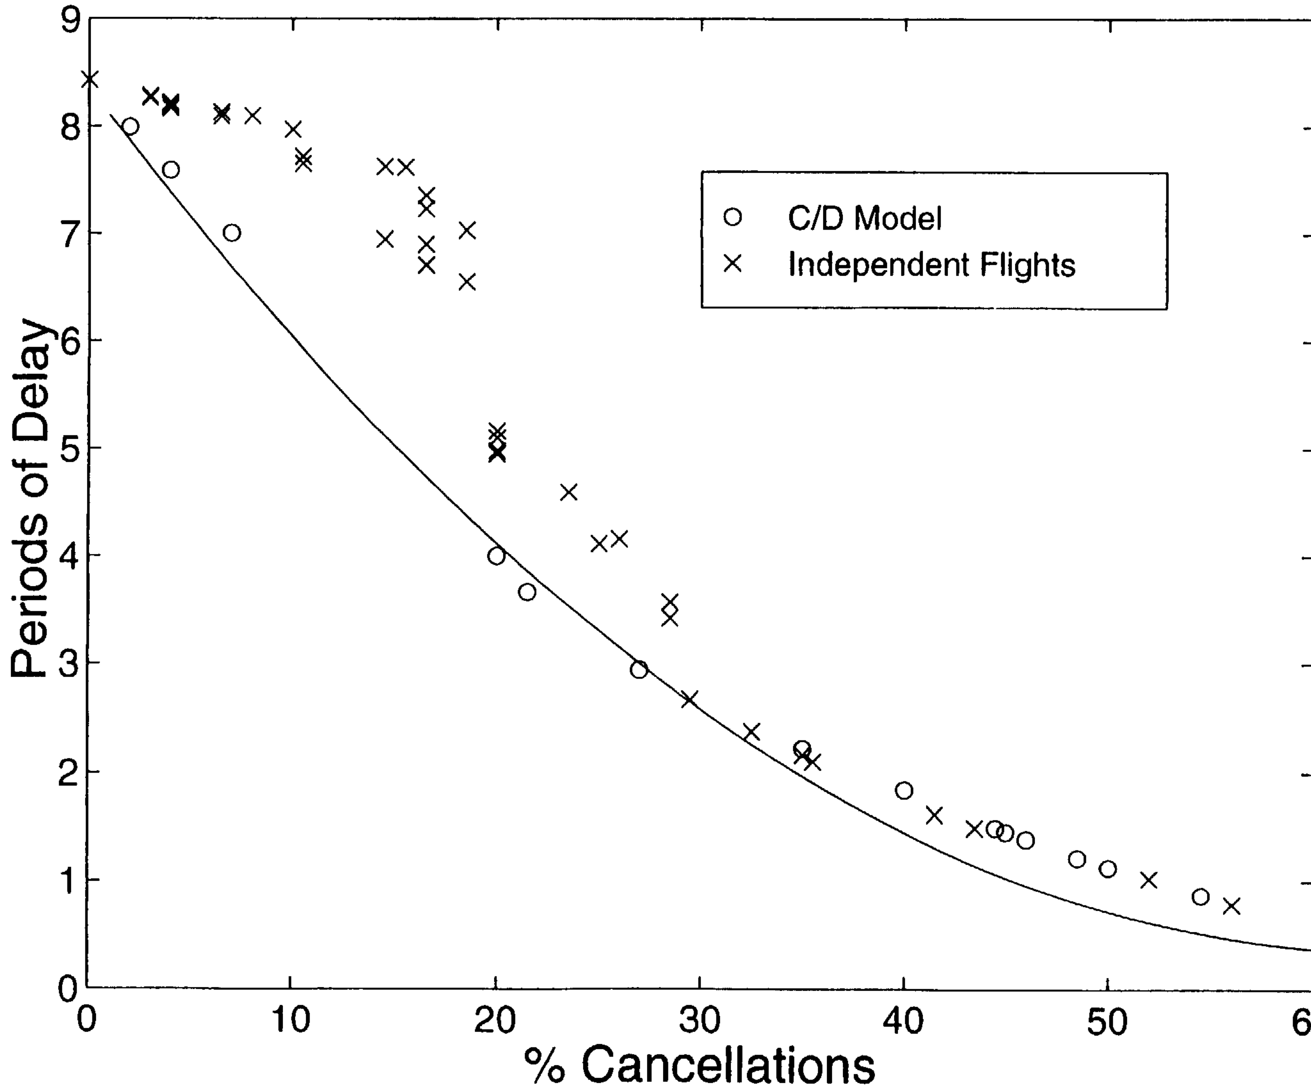

IF = Independent Flights

C/D = Cancellation/Delay Model

Assignment of Flights to Slots by C/D Algorithm

(i) Slot used by a flight from bank i; the flight arrives with the rest of the bank.

△ Slot used by a flight from bank i; the flight is separated from the rest of the bank.

□ Unused slot

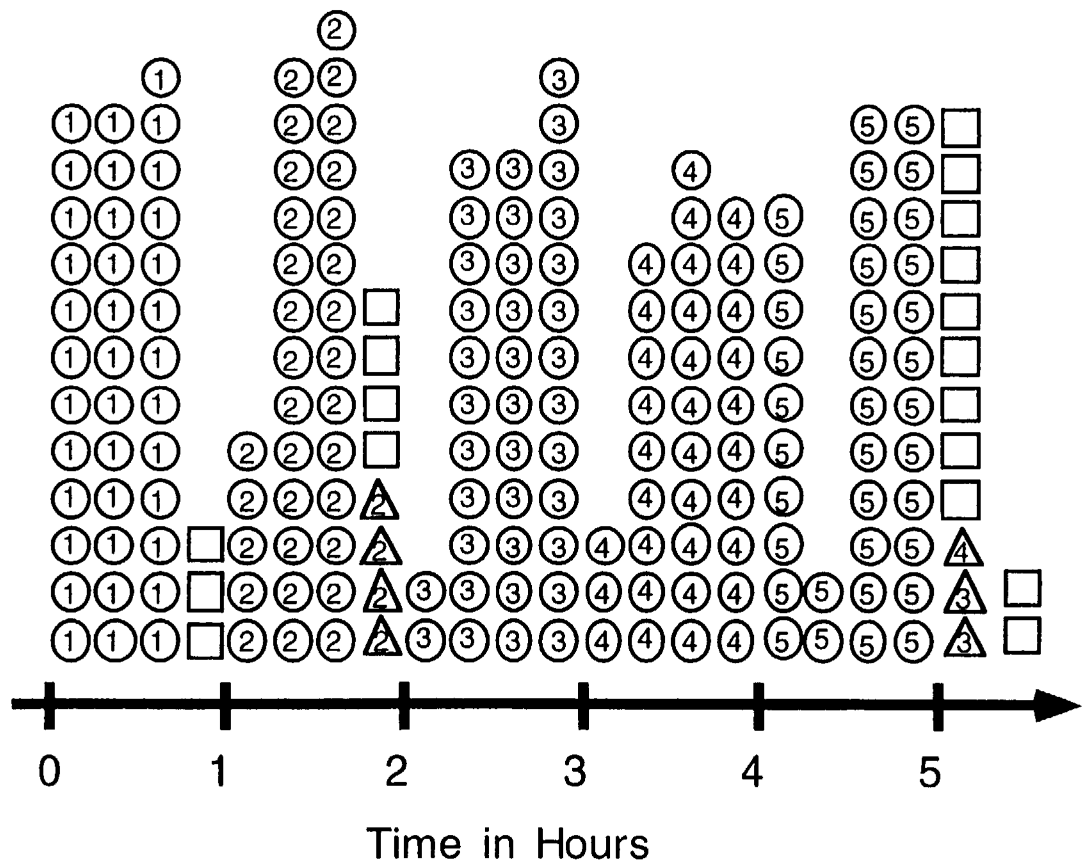

Current Experiments: USAir Data

Bank #	Flights/bank	Bank Start Time	Bank End Time
1	50	7:15 am	8:15 am
2	44	8:45 am	9:30 am
3	23	10:00 am	10:30 am
4	51	11:45 am	12:30 pm
5	35	12:30 pm	1:15 pm
6	43	2:00 pm	3:00 pm
7	43	3:15 pm	4:15 pm
8	52	4:30 pm	5:30 pm
9	63	7:15 pm	8:15 pm
10	53	8:30 pm	9:30 pm

Sample Input File

```
21
15
15
.
.
.
15
15
15
15

4
50 4 22 23 27 30 34 37 42 45 48 53 56 61 66 69 71 75 78 81
44 9 22 26 27 31 33 34 39 41 46 47 50 54 55
23 13 28 33 38 39 40 44 49 52 55
51 21 27

168
1 17 2 7 4 2 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
1 14 4 5 2 1 0 0 0 0 0 0 0 0 0 0 0 0 0 0
1 17 2 5 2 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
1 17 4 6 4 2 1 0 0 0 0 0 0 0 0 0 0 0 0 0
1 17 2 6 5 3 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
1 16 4 6 4 2 0 0 0 0 0 0 0 0 0 0 0 0 0 0
1 16 4 6 5 4 2 0 0 0 0 0 0 0 0 0 0 0 0 0
.
.
.
4 14 20 6
4 18 19 6 3
4 14 20 6
4 17 19 7 5
4 14 19 8 6
4 13 20 5
4 13 19 5 4
4 13 20 8
4 18 20 8
```

Sample Output File

```
BANK 1   COMPLETION TIME: 5  (SCHEDULED: 4)  ***

 Flight number  Arrival Time  Scheduled Arrival Time
 -------------  ------------  ----------------------

    1-1          2            2
    1-2          4            4
    1-3          2            2
...
    3-20          12           12
    3-21          13           13
    3-22          13           13
    3-23          12           12

BANK 4   COMPLETION TIME: 21  (SCHEDULED: 21)

 Flight number  Arrival Time  Scheduled Arrival Time
 -------------  ------------  ----------------------

    4-1          19           19
    4-2          20           20
    4-3          19           19
    4-4          19           19
    4-5        CANCELLED          21  ***
...
    4-44          19           19
    4-45          20           20
    4-46          21           19  ***
    4-47          20           19  ***
    4-48        CANCELLED         20  ***
    4-49        CANCELLED         19  ***
    4-50        CANCELLED         20  ***
    4-51          21           20  ***
```

Cost of solution = 179.000000

Number of iterations = 7388

Percent cancelled = 3.571429

Percent delayed = 4.166667

Number of rows: 8223

Number of columns: 3851

Solution Run-Times

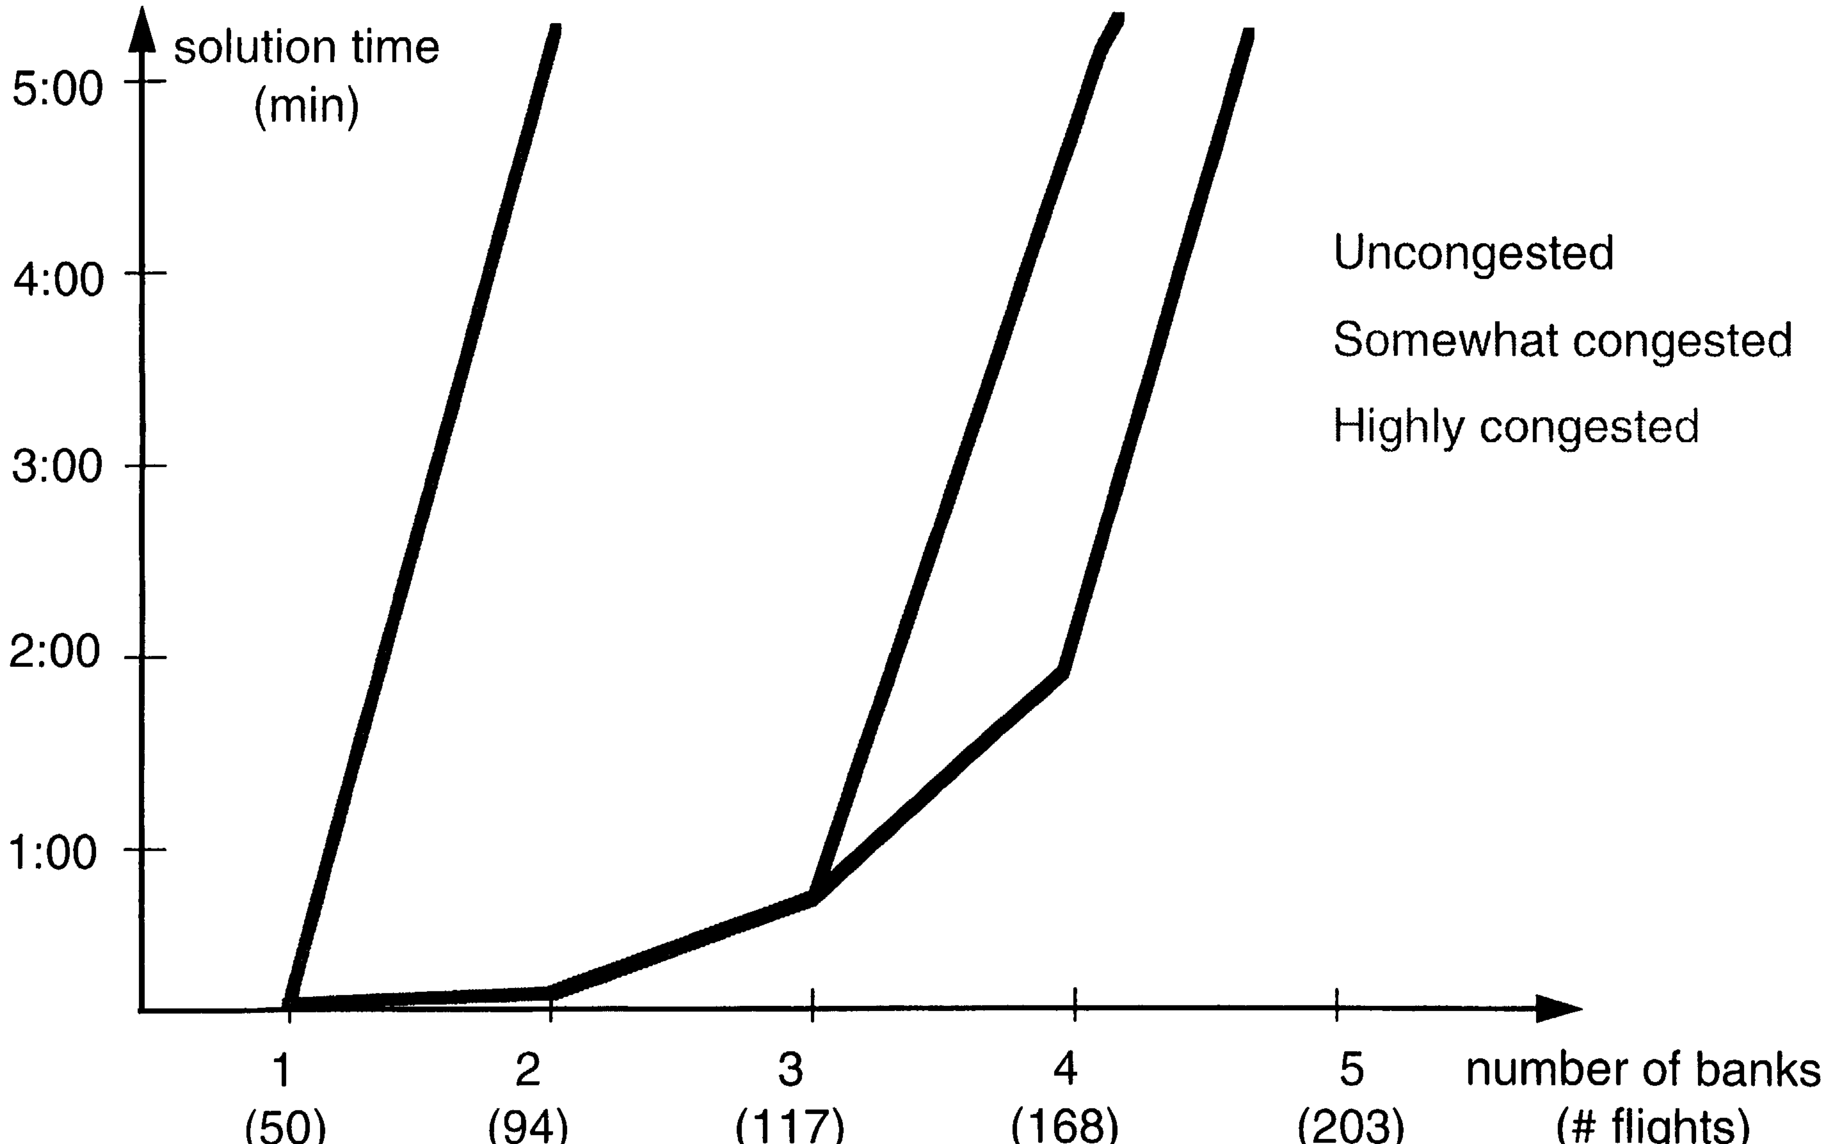

Observations/Restrictions

- Independent flights are modeled as one-flight "banks" (inefficient)
- No penalty for delaying an individual flight (inside its bank)
- No bank overlap is permitted

Example Future Operational Scenario

- **Initial forecast: VMC—no delays**
- **During the day, forecast changes to IMC at one or more airports**
- **The FAA runs a ground delay program in the ATCSCC**
- **Each airline is given a set of landing slots—CTA's**
- **Each airline runs algorithms to reduce delay on its high priority flights—model airline technology by choice of algorithms**
  - Tactical arrival window assignment
  - Tactical maximum marginal return assignment (e.g., ABG)
    - Priorities model airline operational policies
  - Tactical decision-making for slot utilization
- **Each airline's desired sequences at each airport are provided to ARTCCs and TRACONs**
- **CTAS utilizes airline sequence preferences**
- **Performance evaluation based on system-level metrics**

Testbed Environment

Testbed Top Level View

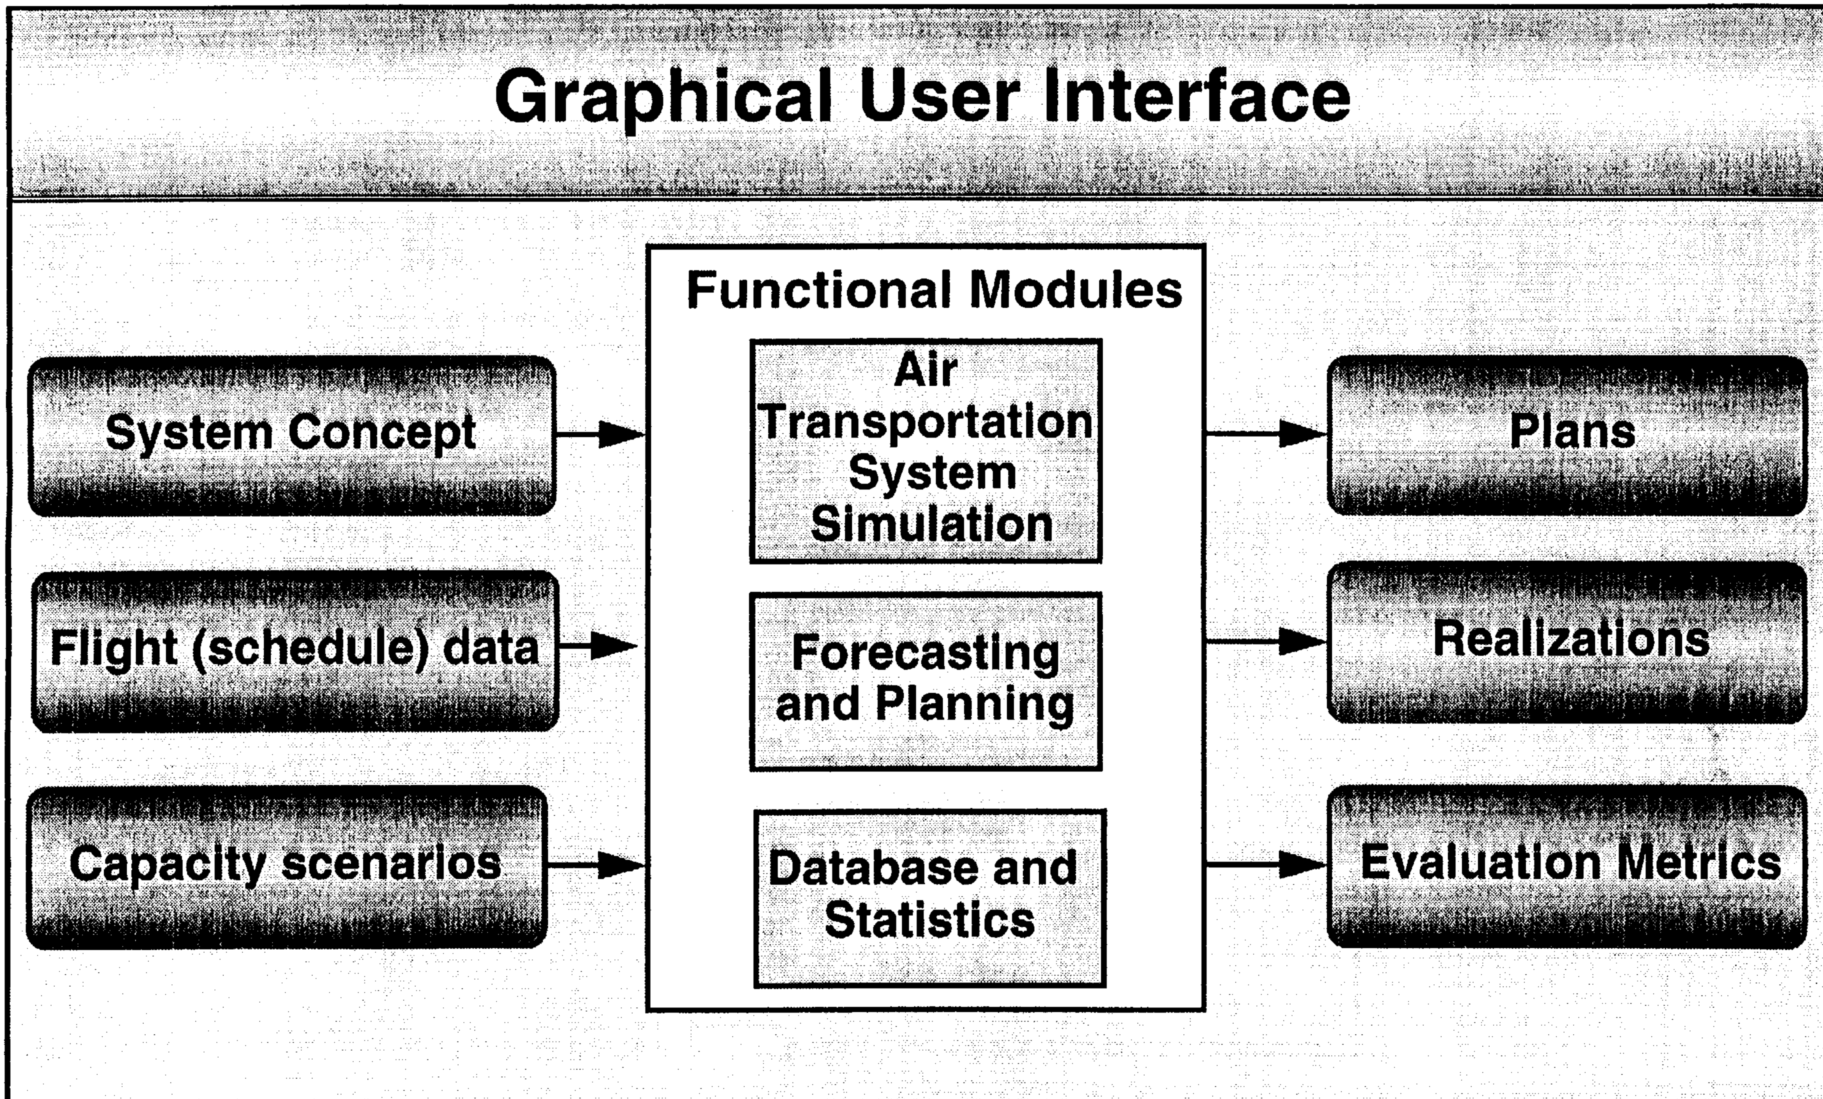

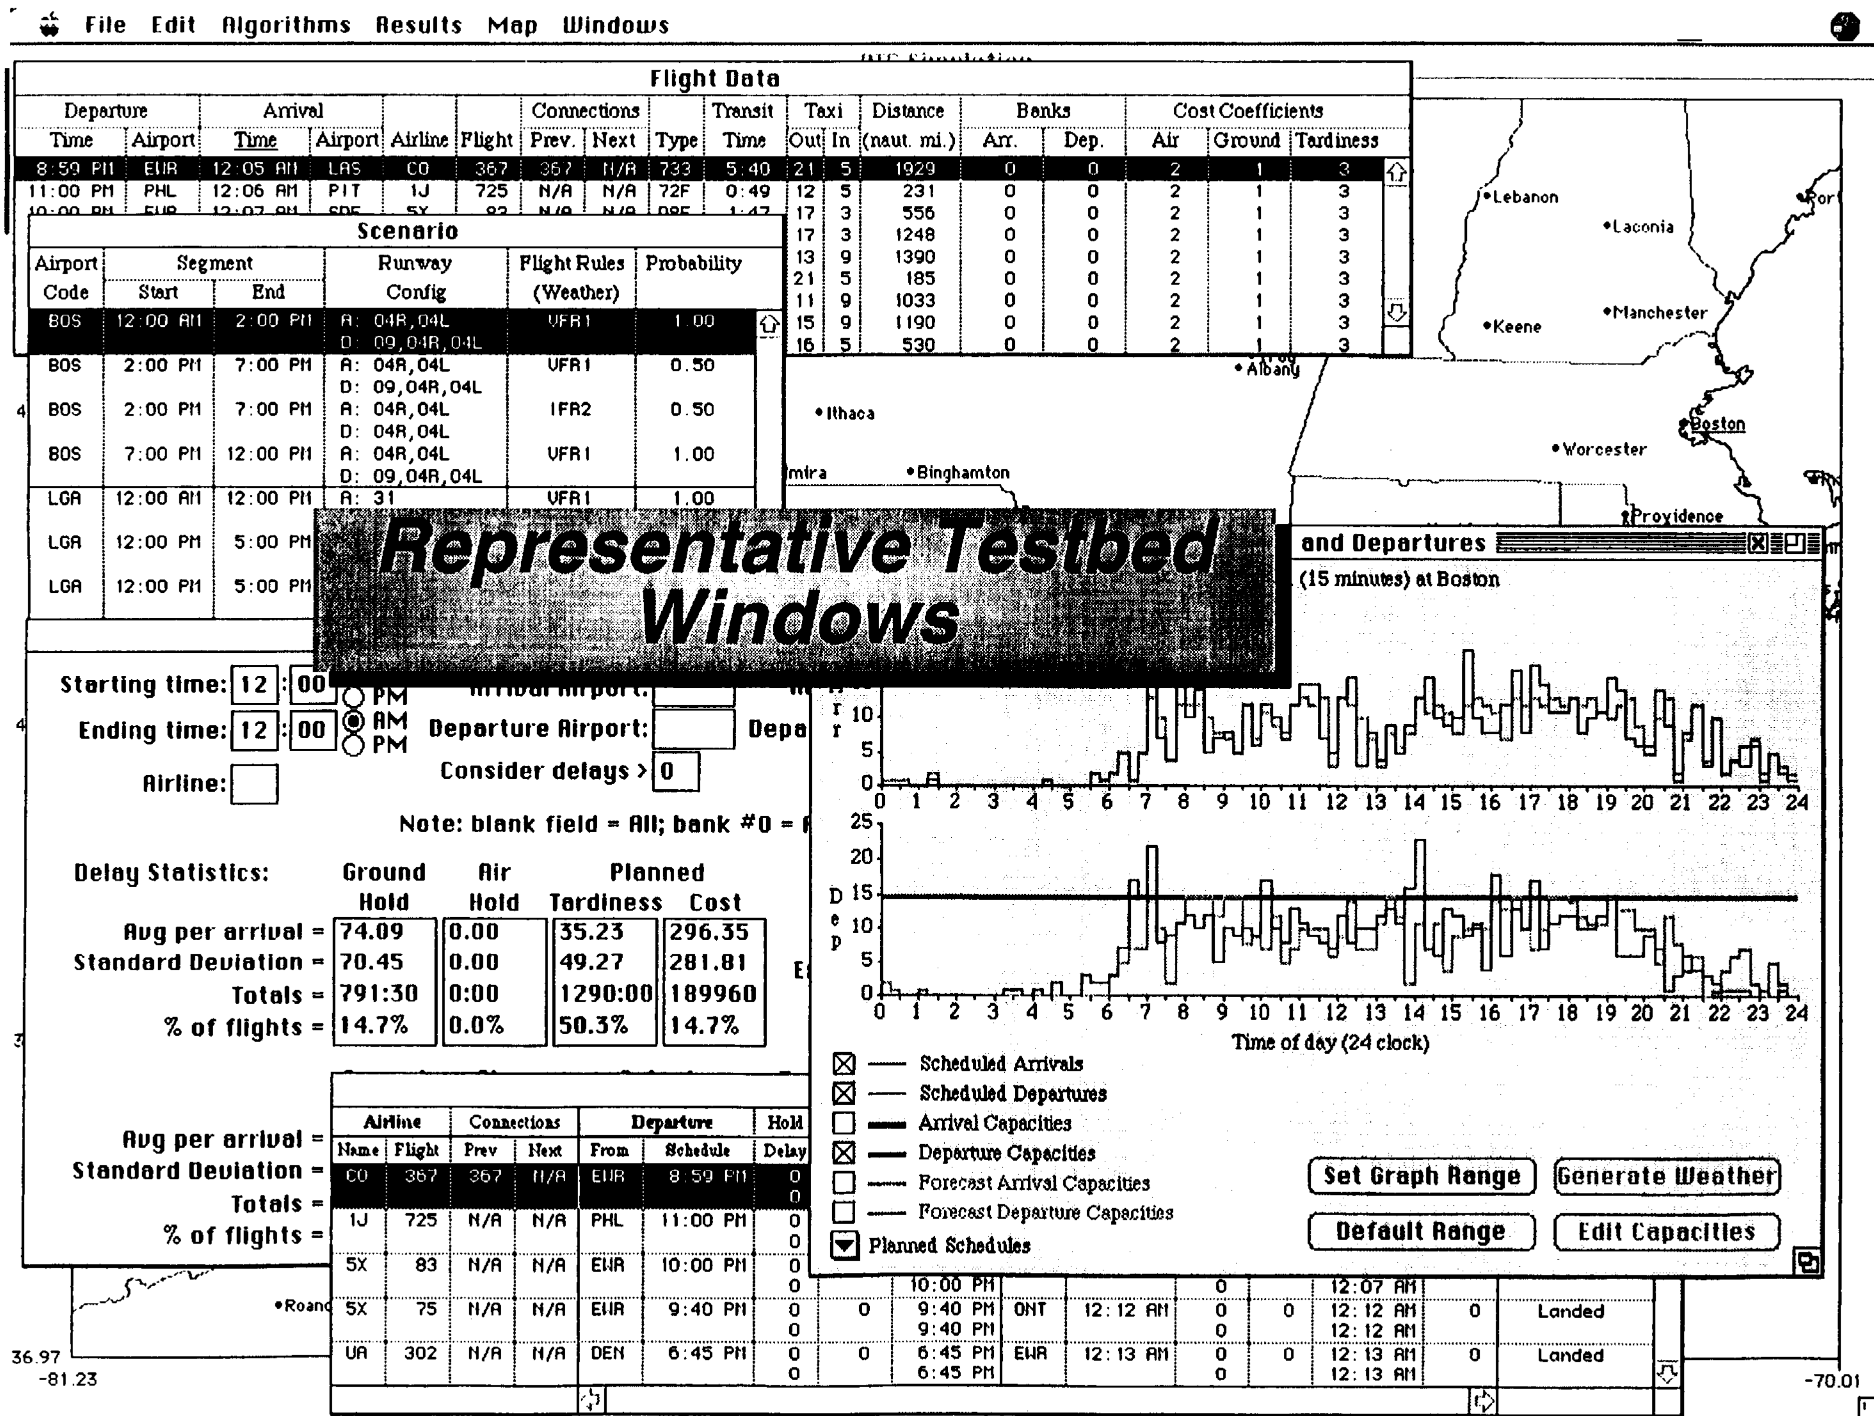
Representative Testbed Windows
File Edit Algorithms Results Map Windows
Flight Data
Scenario
Starting time: 12 : 00
Ending time: 12 : 00
Airline:
Departure Airport:
Consider delays > 0
Note: blank field = All; bank #0 =
Delay Statistics:
Ground Hold
Air Hold
Planned Tardiness
Cost
Avg per arrival =
Standard Deviation =
Totals =
% of flights =
and Departures
(15 minutes) at Boston
Time of day (24 clock)
Scheduled Arrivals
Scheduled Departures
Arrival Capacities
Departure Capacities
Forecast Arrival Capacities
Forecast Departure Capacities
Planned Schedules
Set Graph Range
Generate Weather
Default Range
Edit Capacities
Lebanon
Laconia
Keene
Manchester
Albany
Ithaca
Worcester
Boston
Binghamton
Providence
36.97
-81.23
-70.01

Extensions in Simulation Capabilities

- **Transit times different from scheduled times**
- **Control probability distribution of capacity period by period**
- **Model**
  - Taxi-out
  - Airborne portion
  - Taxi-in
  - Departure capacities
  - Hubs/Banks
  - UPR
    - Reduced transit times in airborne portion of flight
  - CTAS
    - Increase airport acceptance rate
  - SMA
    - Reduce taxi-out delays

Extensions in Planning and Evaluation Capabilities

- **Planning algorithms**
  - Managed Airborne Reservoir (MAR)
  - Regional ground delay programs
  - Optimal ground delay assignment
  - Heuristic ground delay assignment
  - FAA-airlines collaborative slot assignment
    - Tactical arrival window assignment
    - Tactical maximum marginal return assignment
  - Plan to schedule or transit time
- **Metrics**
  - Planned
    - *Gate (=ground) delay*
    - Air delay
    - Tardiness amount and cost
  - Actual
    - *Gate delay*
    - Taxi-out delay
    - *Air delay*
    - Taxi-in delay
    - Tardiness amount and cost

Delay Metrics

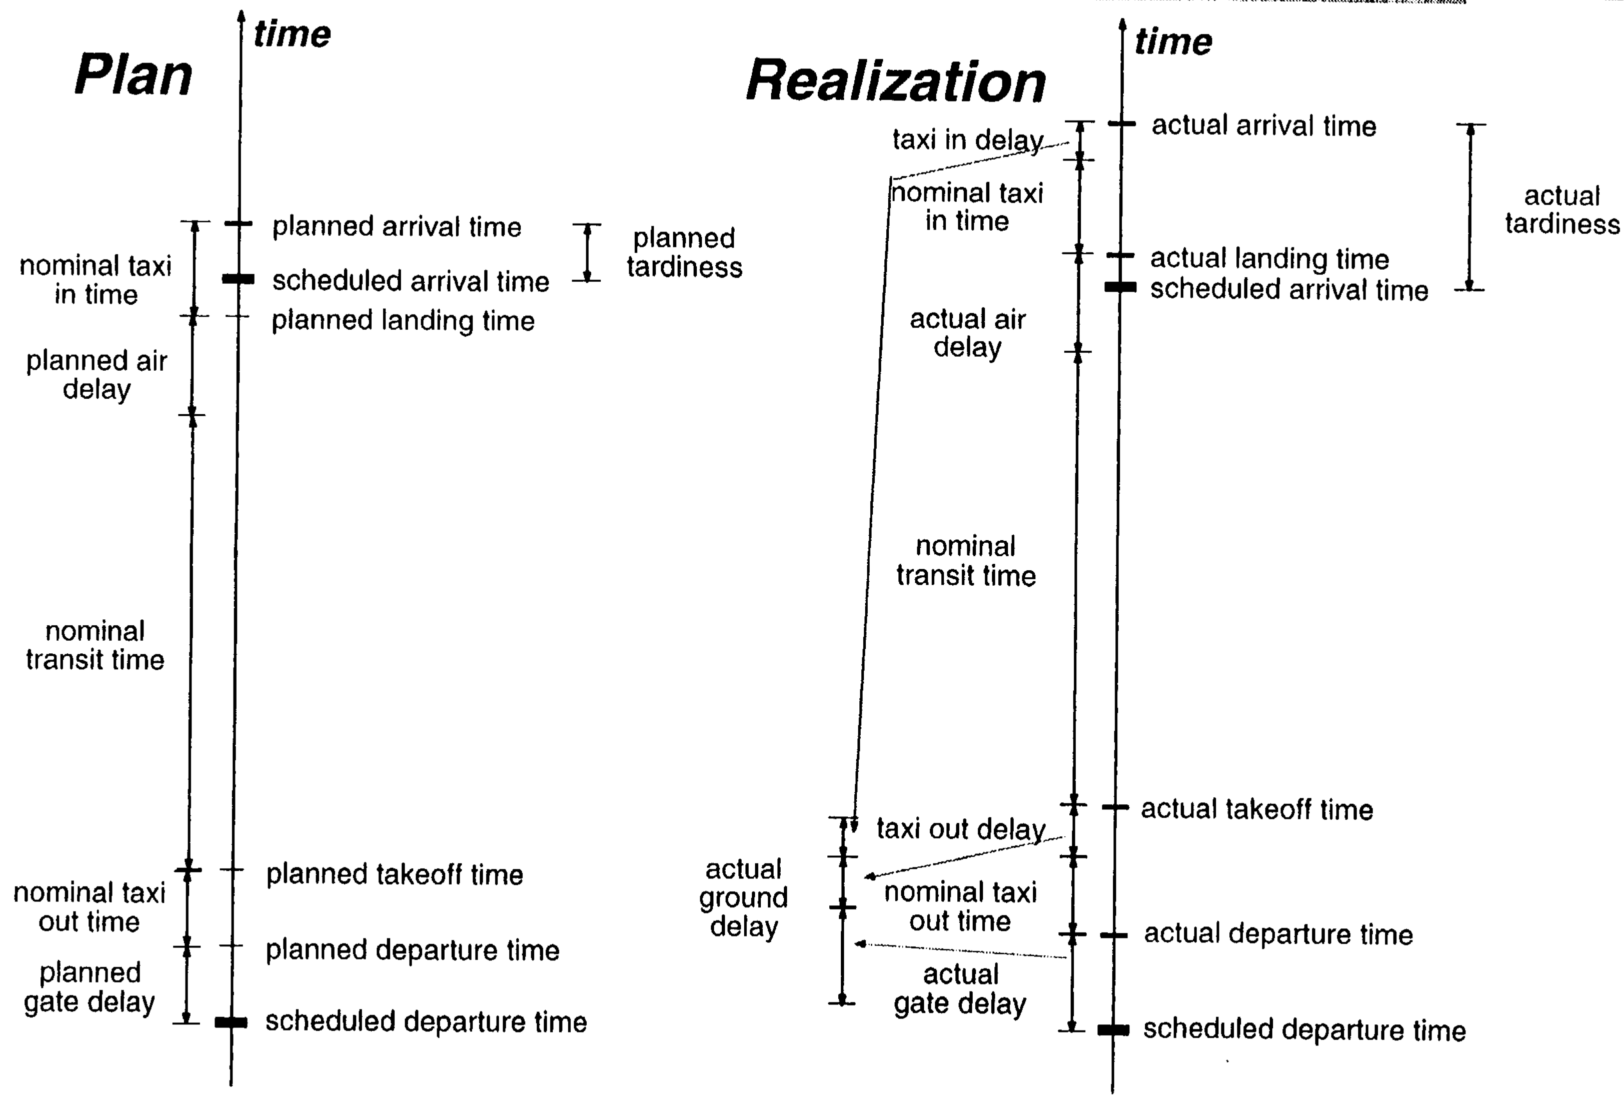

ground delay = gate delay + taxi-out delay + taxi-in delay

Other Models

- **Provide input to Testbed**
  - Schedule generation—POAGG
  - Weather and capacity modeling

- **Not integrated yet**
  - Higher fidelity enroute and transition area framework
  - Modeling airline slot utilization
  - Slot allocation and CTA specification under uncertainty
  - Real-time enroute rerouting

Data Needs

Actual and Planned Flight Data

Six months of daily records

1. ETMS or other data (such as SAR) showing last flight plan before departure, actual position updates, and departure and arrival times
2. Counts of aircraft in sectors, in 5 or 10-minute intervals
3. Estimates of potential conflicts for individual aircraft flights
4. For OAG flights, scheduled and actual aircraft itineraries; bank assignments (for these data, representative days sprinkled over the past few years are sufficient)
5. For GA traffic flying VFR, radar position reports.
6. For GA and military: filed flight plans versus flown plans (note that GA often changes filed plans in mid-flight and these may or may not be archived)
7. GA position reports for those aircraft flying VFR.

Transit times

1. Minimum, planned (or expected), and actual takeoff to landing times, by flight segment and by aircraft
2. Minimum, expected and actual taxi-out and taxi-in times
3. Time spent on ramp away from gate

Weather

- **Historical actual and forecasted airport surface weather (1985-1995, or similar multi-year period)**
  1. Wind speed and direction
  2. Ceiling
  3. Visibility
- **Historical actual and forecasted high altitude weather (1985-1995, or similar multi-year period)**
  1. Wind (3-D gridded values)
  2. Storms (how are these characterized? height, degree of lightning, turbulence, intensity of precipitation or some other features that are more directly related to sector capacity?)

Capacities

- **Airport capacity for each pacing airport**
  1. List and description of possible runway configurations
  2. Observed configurations for the last few years, correlated with the weather observations from I above
  3. The ceiling and visibility values that separate the weather conditions (e.g., VFR1, VFR2, IFR1, IFR2), by configuration if necessary
  4. Arrival and departure capacities for each runway configuration and weather condition
  5. Magnetic deviation
- **Sector boundaries and capacities**
  1. Sector boundaries
  2. Sector capacities, and the FAA algorithm used to calculate them
  3. Sector capacities as a function of high altitude weather conditions
  4. Airways coordinates
  5. Boundaries of special-use airspace
  6. Descriptions of reasons for SUA restrictions, times restrictions are in effect, and organizations responsible for approving relief from restrictions (if any)

Summary

- Our project's primary objective: to perform system-level analyses of collaborative tactical flow management under Free Flight
- The Draper ATFM Testbed provides the integrating environment